



Carnegie Mellon
Software Engineering Institute

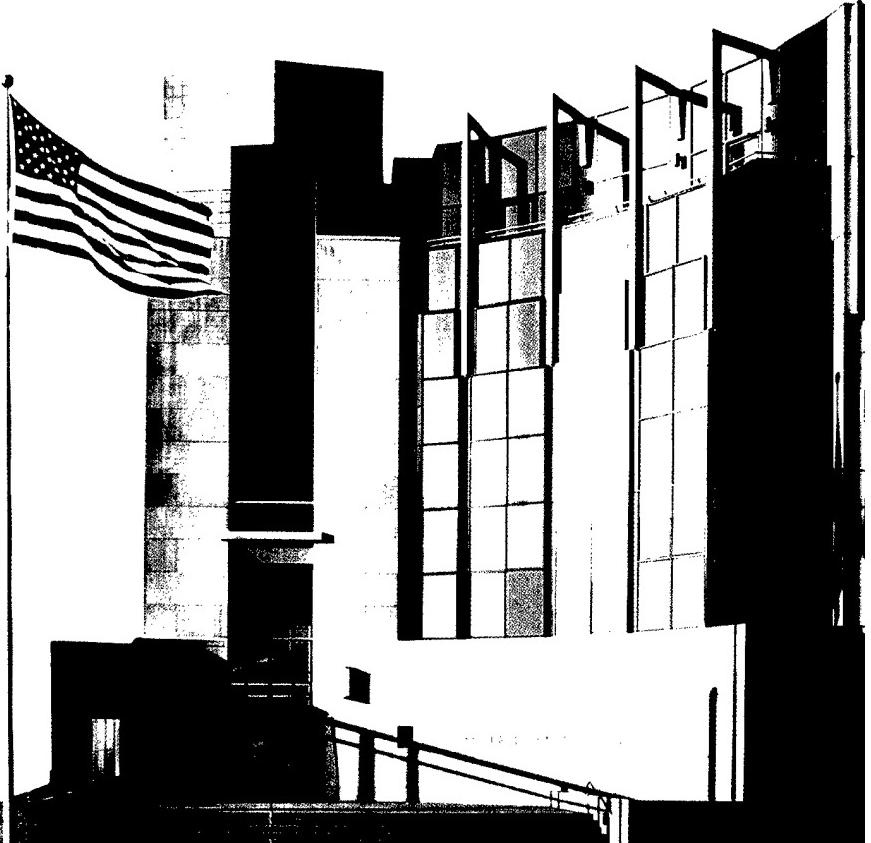
Detecting Signs of Intrusion

Julia Allen
Ed Stoner

October 2000

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

SECURITY IMPROVEMENT MODULE
CMU/SEI-SIM-009



20001107 031

Carnegie Mellon University does not discriminate and Carnegie Mellon University is required not to discriminate in admission, employment, or administration of its programs or activities on the basis of race, color, national origin, sex or handicap in violation of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972 and Section 504 of the Rehabilitation Act of 1973 or other federal, state, or local laws or executive orders.

In addition, Carnegie Mellon University does not discriminate in admission, employment or administration of its programs on the basis of religion, creed, ancestry, belief, age, veteran status, sexual orientation or in violation of federal, state, or local laws or executive orders. However, in the judgment of the Carnegie Mellon Human Relations Commission, the Department of Defense policy of "Don't ask, don't tell, don't pursue" excludes openly gay, lesbian and bisexual students from receiving ROTC scholarships or serving in the military. Nevertheless, all ROTC classes at Carnegie Mellon University are available to all students.

Inquiries concerning application of these statements should be directed to the Provost, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-6684 or the Vice President for Enrollment, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-2056.

Obtain general information about Carnegie Mellon University by calling (412) 268-2000.



CarnegieMellon
Software Engineering Institute

Pittsburgh, PA 15213-3890

Detecting Signs of Intrusion

CMU/SEI-SIM-009

Julia Allen

Ed Stoner

October 2000

Networked Systems Survivability Program

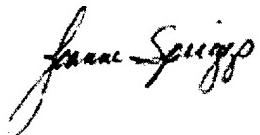
Unlimited distribution subject to the copyright.

This report was prepared for the

SEI Joint Program Office
HQ ESC/AXS
5 Eglin Street
Hanscom AFB, MA 01731-2116

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

FOR THE COMMANDER



Joanne E. Spriggs
Contracting Office Representative

Updates to this report and the effort to produce it were sponsored by the US Air Force Computer Resources Support Improvement Program (CRSIP) in collaboration with the Air Force Information Warfare Center (AFIWC). The original versions of this report (1997—Detecting Signs of Intrusion, 1998—Preparing to Detect Signs of Intrusion) and the effort to produce them were sponsored by the SEI primary sponsor.

The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

“CERT” and “CERT Coordination Center” are registered in the U.S. Patent and Trademark Office.

Copyright © 2000 by Carnegie Mellon University.

Requests for permission to reproduce this document or to prepare derivative works of this document should be addressed to the SEI Licensing Agent.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This work was created in the performance of Federal Government Contract Number F19628-95-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 52.227-7013.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (<http://www.sei.cmu.edu/publications/pubweb.html>).

Table of Contents

Preface	iii
Detecting Signs of Intrusion	1
1. Establish a policy and procedures that prepare your organization to detect signs of intrusion.	9
2. Identify data that characterize systems and aid in detecting signs of suspicious behavior.	15
3. Manage logging and other data collection mechanisms.	27
4. Ensure that the software used to examine systems has not been compromised.	31
5. Monitor and inspect network activities for unexpected behavior.	35
6. Monitor and inspect system activities for unexpected behavior.	41
7. Inspect files and directories for unexpected changes.	49
8. Investigate unauthorized hardware attached to your organization's network.	53
9. Inspect physical resources for signs of unauthorized access.	55
10. Review reports by users and external contacts about suspicious and unexpected behavior.	57
11. Take appropriate actions upon discovering unauthorized, unexpected, or suspicious activity.	61

Preface

This document is one of a series of publications of the Software Engineering Institute at Carnegie Mellon University called *security improvement modules*. They are intended to provide practical guidance to help organizations improve the security of their networked computer systems.

Module structure	Each module addresses an important but relatively narrowly defined problem in network and system security. The first section of the module describes the problem and outlines a set of <i>security improvement practices</i> to help solve it. Each practice is a recommended way of performing common tasks related to the secure operation of networked computer systems. The remaining sections of the module are detailed descriptions of the practices. Each includes a rationale for the recommended actions and a description of how to perform them.
Intended audience	The practices are primarily written for system and network administrators whose day-to-day activities include installation, configuration, and maintenance of the computers and networks. Occasionally, practices are written to assist the managers responsible for network and system administration.
Revised versions	Network and system technologies continue to evolve rapidly, leading to new security problems and solutions. Modules and practices need to be revised occasionally, so to permit more timely publication of new versions, we also publish them on the World Wide Web. At the end of each section of this document is the URL of its web version.
Implementation details	How an organization adopts and implements the practices often depends on the networking and computing technologies it uses. For some practices, technology-specific implementation details are published on the World Wide Web. The Web version of each practice contains links to the implementation details.

Acknowledgments

Updates to this report and the effort to produce it were sponsored by the US Air Force Computer Resources Support Improvement Program (CRSIP) in collaboration with the Air Force Information Warfare Center (AFIWC). The original versions of this report (1997—Detecting Signs of Intrusion, 1998—Preparing to Detect Signs of Intrusion) and the effort to produce them were sponsored by the SEI primary sponsor.

The authors acknowledge contributions made to this report by the authors of version 1, all of whom are affiliated with the SEI unless otherwise noted:

Detecting Signs of Intrusion, 1997

- Robert Firth
- Gary Ford
- Barbara Fraser
- John Kochmar
- Suresh Konda
- John Richael
- Derek Simmel
- Lisa Cunningham, Computer Sciences Corporation

Preparing to Detect Signs of Intrusion, 1998

- John Kochmar
- Christopher Alberts
- Cory Cohen
- Gary Ford
- Barbara Fraser
- Suresh Konda
- Klaus-Peter Kossakowski
- Derek Simmel

and by the reviewers of this report:

- Audrey Dorofee
- Chad Dougherty
- Eric Hayes
- Klaus-Peter Kossakowski
- Barbara Laswell
- Rudy Maceyko
- Jerome Marella
- Larry Rogers
- Bradford Willke

Detecting Signs of Intrusion¹

Intruders are always looking for new ways to break into networked computer systems. They may attempt to breach your network's perimeter defenses from remote locations, or try to physically infiltrate your organization to access information resources. Intruders seek old, unpatched vulnerabilities as well as newly discovered vulnerabilities in operating systems, network services, and protocols and take advantage of both. They develop and use sophisticated programs to rapidly penetrate systems. As a result, intrusions and the damage they cause can be achieved in seconds.²

Even if your organization has implemented comprehensive information security protection measures (such as firewalls), it is essential that you closely monitor your information assets and transactions involving these assets for signs of intrusion. Monitoring may be complicated because intruders often hide their activities by changing the systems they break into. An intrusion may have already happened without you noticing because everything *seemed* to be operating normally.

A general security goal is to prevent intrusions. However, because no prevention measures are perfect, you also need a strategy for handling intrusions that includes *preparation*, *detection*, and *response*. This module focuses on preparation and detection. The practices recommended below are designed to help you prepare for and detect intrusions by looking for unexpected or suspicious behavior and “fingerprints” of known intrusion methods.

Who should read these practices

These practices are intended primarily for system and network administrators, managers of information systems, and security personnel responsible for networked information resources.

These practices are applicable to your organization if its networked systems infrastructure includes

- host systems providing services to multiple users (file servers, timesharing systems, database servers, web servers, etc.)

1. This module replaces the previous versions of *Detecting Signs of Intrusion* [1997] and *Preparing to Detect Signs of Intrusion* [1998]. We have added information about asset characterization and system and network monitoring.

2. Refer to Figure 1-2 and the accompanying description in *State of the Practice of Intrusion Detection Technologies* [Allen 00].

- local-area or wide-area networks
 - direct connections, gateways, or modem access to and from external networks, such as the Internet
-

What these practices do not cover	<p>These practices do not cover</p> <ul style="list-style-type: none">• preventing intrusions• responding to intrusions. Refer to <i>Responding to Intrusions</i> [Kossakowski 99].• establishing initial configurations of applications, operating systems, networks, or workstations. Refer to <i>Securing Desktop Workstations</i> [Ford 99], <i>Securing Network Servers</i> [Allen 00], and <i>Securing Public Web Servers</i> [Kossakowski 00].• protecting user privacy while in the process of detecting signs of intrusion• using security monitoring and reporting services provided by outside (third party) organizations
--	---

Security issues	<p>If you do not know that an intrusion or an intrusion attempt has occurred, it is difficult, if not impossible, to later determine if your systems have been compromised. If the information necessary to detect an intrusion is not being collected and reviewed, you cannot determine what sensitive data, systems, and networks are being attacked and what breaches in confidentiality, integrity, or availability have occurred. As a result of an inadequate ability to detect signs of intrusion, the following may occur:</p> <ol style="list-style-type: none">1. You will be unable to detect such signs in a timely manner due to the absence of necessary warning mechanisms and review procedures.2. You will be unable to identify intrusions because of the absence of expected state information with which to compare your current operational state. Differences between this expected configuration and your current state can provide an indication that an intrusion has occurred.3. You will be unable to determine the full extent of the intrusion and the damage it has caused. You will also be unable to tell whether or not you have completely removed the intruder from your systems and networks. This will significantly increase your time to recover.4. Your organization may be subject to legal action. Intruders make use of systems they have compromised to launch attacks against others. If one of your systems is used in this way, you may be held liable for not exercising adequate due care with respect to security.5. Your organization may experience lost business opportunities and its reputation may suffer.
------------------------	---

If you are adequately prepared and if you have the necessary policies and procedures in place to detect signs of intrusion, then you can mitigate your risk of exposure to intrusion and mitigate possible damage to your systems.

Security improvement approach	<p>These practices assume that</p> <ul style="list-style-type: none">• You have performed security planning (such as policy formulation, disaster recovery and business continuity planning, risk assessment, identification of critical information assets) that addresses your organization's business objectives.
--------------------------------------	--

- You have performed trade-off analyses to determine the cost of protecting versus the cost of reconstituting critical assets (data, systems, networks, workstations, tools) in the event of an intrusion. Protecting an asset includes consideration of the loss of confidentiality and customer confidence if the asset is disclosed (e.g., confidential, competitive information). It is likely not feasible to protect all assets.
- You have a documented disaster recovery policy and procedures that include determining what assets are critical to protect and with what priority. The policy identifies who has responsibility for and authority to access each asset that needs to be recovered, under what conditions, and by what means.

The general approach to detecting intrusions is

1. Observe your systems for anything unexpected or suspicious.
2. Investigate anything you find to be unusual.
3. If your investigation finds something that isn't explained by authorized activity, immediately initiate your intrusion response procedures.

While this process sounds simple enough, implementing it is a resource-intensive activity that requires continuous, automated support and daily administrative effort. Furthermore, the scale of intrusion detection practices may need to change as threats, system configurations, or security requirements change. In all cases, however, there are five areas that must be addressed:

- adequate preparation, which should include defining the required policies and procedures necessary to meet your business objectives and prepare your staff and systems to detect signs of intrusion
- integrity of the software you use to detect intrusions
- monitoring the behavior of your systems and the traffic on your networks
- physical forms of intrusion to your computer systems, offline data storage media, and output devices
- follow through, including investigation of reports by users and other reliable sources (such as incident response teams) and taking action when unexpected activities occur

As you look for signs of intrusion, keep in mind that information from one source may not appear suspicious by itself. Inconsistencies among several sources can sometimes be the best indication of suspicious behavior or intrusions.

Summary of recommended practices	Area	Recommended Practice
	Preparation	<ol style="list-style-type: none"> 1. Establish a policy and procedures that prepare your organization to detect signs of intrusion. 2. Identify data that characterize systems and aid in detecting signs of suspicious behavior. 3. Manage logging and other data collection mechanisms.
	Integrity of intrusion detection software	<ol style="list-style-type: none"> 4. Ensure that the software used to examine systems has not been compromised.
	Behavior of networks and systems	<ol style="list-style-type: none"> 5. Monitor and inspect network activities for unexpected behavior. 6. Monitor and inspect system activities for unexpected behavior. 7. Inspect files and directories for unexpected changes.
	Physical forms of intrusion	<ol style="list-style-type: none"> 8. Investigate unauthorized hardware attached to your organization's network. 9. Inspect physical resources for signs of unauthorized access.
	Follow through	<ol style="list-style-type: none"> 10. Review reports by users and external contacts about suspicious and unexpected behavior. 11. Take appropriate actions upon discovering unauthorized, unexpected, or suspicious activity.

Abbreviations used in these practices		
	ACK	Acknowledgement
	ARP	Address Resolution Protocol
	ASCII	American Standard Code for Information Interchange
	ACL	Access Control List
	BOOTP	Boot Protocol
	CGI	Common Gateway Interface
	CPU	Central Processing Unit
	DHCP	Dynamic Host Configuration Protocol
	DNS	Domain Name System
	FTP	File Transfer Protocol
	HAVAL	Hashing Algorithm with VAriable Length
	HTTP	Hypertext Transfer Protocol
	ICMP	Internet Control Message Protocol
	IDS	Intrusion Detection System
	IP	Internet Protocol
	IRC	Internet Relay Chat
	ISP	Internet Service Provider
	MAC	Media Access Control
	NTP	Network Time Protocol
	PGP	Pretty Good Privacy
	SHA	Secure Hash Algorithm
	SNMP	Simple Network Management Protocol

SYN	Synchronize
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
WORM	Write Once Read Many

References and sources	[Alberts 00]	Alberts, Christopher J., et al. "Health Information Risk Assessment and Management: Toolkit Section 4.5." <i>CPRI Toolkit: Managing Information Security in Health Care, Version 2</i> [online]. Available WWW: <URL: http://www.3com.com/healthcare/securitynet/hipaa/4_5.html > (2000).
	[Alberts 99]	Alberts, Christopher J., et al. <i>Operationally Critical Threat, Assets, and Vulnerability EvaluationSM(OCTAVESM) Framework, Version 1.0.</i> (CMU/SEI-99-TR-017). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1999 [online]. Available WWW: <URL: http://www.sei.cmu.edu/publications/documents/99.reports/99tr017/99tr017abstract.html >.
	[Allen 99]	Allen, Julia, et al. <i>State of the Practice of Intrusion Detection Technologies.</i> (CMU/SEI-99/TR-028). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1999 [online]. Available WWW: <URL: http://www.sei.cmu.edu/publications/documents/99.reports/99tr028/99tr028abstract.html >.
	[Allen 00]	Allen, Julia, et al. <i>Securing Network Servers.</i> (CMU/SEI-SIM-010). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2000 [online]. Available WWW: <URL: http://www.cert.org/security-improvement/modules/m10.html >.
	[Bejtlich]	Bejtlich, Richard. <i>Interpreting Network Traffic: A Network Intrusion Detector's Look at Suspicious Events</i> [online]. Available WWW: <URL: http://packetstorm.security.com/papers/IDS/intv2.txt > (1999).
	[CERIAS]	Center for Education, Research, and Information Assurance Security (CERIAS) [formerly known as Computer Operations, Audit, and Security (COAST)]. Monitoring and intrusion detection tools available for downloading [online]. Available WWW: <URL: http://www.cerias.purdue.edu > (2000).
	[CERT/CC]	CERT® Coordination Center. Advisories, incident notes, vulnerability notes, and tech tips [online]. Available WWW: <URL: http://www.cert.org >. Relevant tech tips include <i>Intrusion Detection Checklist</i> [online] (available WWW: <URL: http://www.cert.org/tech_tips/intruder_detection_checklist.html >) and <i>Steps for Recovering from a UNIX Root Compromise</i> [online] (available WWW: <URL: http://www.cert.org/tech_tips/root_compromise.html >) (2000).
	[Dunigan 99]	Dunigan, Tom & Hinkel, Greg. "Intrusion Detection and Intrusion Prevention on a Large Network: A Case Study." <i>Proceedings of the 1st Workshop on Intrusion Detection and Network Monitoring.</i> Santa Clara, CA. April 9-12, 1999 [online]. Available WWW: <URL: http://www.usenix.org/publications/library/proceedings/detection99/full_papers/dunigan/dunigan_html/index.html >.

- [Firth 97] Firth, Robert, et al. *An Approach for Selecting and Specifying Tools for Information Survivability*. (CMU/SEI-97-TR-009). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1997 [online]. Available WWW: <URL: <http://www.sei.cmu.edu/publications/documents/97.reports/97tr009/97tr009abstract.html>>.
- [Garfinkel 96] Garfinkel, S. & Spafford, G. *Practical UNIX and Internet Security, Second Edition*. Sebastopol, CA: O'Reilly & Associates, Inc., 1996.
- [Guttman 97] Guttman, B. & Bagwill, R. *Internet Security Policy: A Technical Guide —Draft*. Gaithersburg, MD: NIST Special Publication 800-XX, 1997 [online]. Available WWW: <URL: <http://csrc.nist.gov/isptg/html/>>.
- [IETF 97] Internet Engineering Task Force Network Working Group. *RFC 2196 Site Security Handbook* [online]. Edited by Barbara Fraser. Available WWW: <URL: <http://www.ietf.org/rfc/rfc2196.txt>> (1997).
- [Kessler 00] Kessler, Gary C. "Web of Worries." *Information Security* (April 2000) [online]. Available WWW: <URL: <http://www.infosecuritymag.com/>>.
- [Kossakowski 99] Kossakowski, Peter, et al. *Responding to Intrusions*. (CMU/SEI-SIM-006). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1999 [online]. Available WWW: <URL: <http://www.cert.org/security-improvement/modules/m06.html>>.
- [Kossakowski 00] Kossakowski, Peter, et al. *Securing Public Web Servers*. (CMU/SEI-SIM-010). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2000 [online]. Available WWW: <URL: <http://www.cert.org/security-improvement/modules/m11.html>>.
- [Maximum 97] Anonymous. *Maximum Security: A Hacker's Guide to Protecting Your Internet Site and Network*. Indianapolis, IN: Sams.net Publishing, 1997 [online]. Available WWW: <URL: http://mx.nsu.ru/Max_Security/>.
- [Newsham 98] Newsham, Tim & Ptacek, Tom. *Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection* [online]. Available WWW: <URL: http://www.nai.com/media/ps/nai_labs/ids.ps> (1998).
- [Northcutt 99] Northcutt, Stephen. *Network Intrusion Detection: An Analyst's Handbook*. Indianapolis, IN: New Rider, 1999.
- [Pichnarczyk 94] Pichnarczyk, Karyn; Weeber, Steve; & Feingold, Richard. *Unix Incident Guide: How to Detect an Intrusion*. (CIAC-2305 R.1). Livermore, CA: Lawrence Livermore National Laboratory, Department of Energy Computer Incident Advisory Capability, December 1994 [online]. Available WWW: <URL: http://ciac.llnl.gov/ciac/documents/CIAC-2305_UNIX_Incident_Guide_How_to_Detect_an_Intrusion.pdf>.
- [Ranum 99] Ranum, Marcus. "Some Tips on Network Forensics." *Computer Security Institute*, 198 (September 1999): 1-8.
- [Reavis 99] Reavis, Jim. "Do you have an intrusion detection response plan?" *Network World Fusion* (September 13, 1999) [online]. Available WWW: <URL: <http://www.nwfusion.com/newsletters/sec/0913sec1.html>>.
- [Ruiu 99] Ruiu, Dragos. *Cautionary Tales: Stealth Coordinated Attack HOWTO* [online]. Available WWW: <URL: http://www.nswc.navy.mil/ISSEC/CID/Stealth_Coordinated_Attack.html> (1999).
- [SANS 00] The SANS Institute. *How To Eliminate The Ten Most Critical Internet Security Threats: The Experts' Consensus, Version 1.25* [online]. Available WWW: <URL: <http://www.sans.org/topten.htm>> (2000).

- [Seifried 00] Seifried, Kurt. "Creating and Preventing Backdoors in UNIX Systems." *SecurityPortal* (June 28, 2000) [online]. Available WWW: <URL: <http://www.securityportal.com/closet/closet20000628.html>>.
- [Sellens 00] Sellens, John. "System and Network Monitoring."; *login*: 25, 3 (June 2000).
- [Simmel 99] Simmel, Derek, et al. *Securing Desktop Workstations*. (CMU/SEI-SIM-004). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1999 [online]. Available WWW: <URL: <http://www.cert.org/security-improvement/m04.html>>.
- [Stevens 94] Stevens, W. Richard. *TCP/IP Illustrated, Volume 1: The Protocols*. Reading, MA: Addison-Wesley, 1994.
- [Summers 97] Summers, Rita C. *Secure Computing*. New York, NY: McGraw-Hill, 1997.

Where to find updates

The latest version of this module is available on the web at URL
<http://www.cert.org/security-improvement/modules/m09.html>

1

Establish a policy and procedures that prepare your organization to detect signs of intrusion.

A security policy defines the rules that regulate how your organization manages and protects its information and computing resources to achieve security objectives. One of the policy's primary purposes in detecting signs of intrusion is to document important information assets¹ and the threats² to those assets that your organization chooses to address.³

Preparation procedures include the actions necessary to observe systems and networks for signs of unexpected behavior, including intrusion. Observation can take the form of monitoring, inspecting, and auditing.⁴ From these procedures, all concerned parties are able to determine the operational steps they need to take to comply with your policy. These steps will thereby uphold the security of your organization's information and networked systems.

Security policies and procedures that are documented, well known, and visibly enforced establish expected user behavior and serve to inform users of their obligations for protecting computing assets. Users include all those who access, administer, and manage your systems and have authorized accounts on your systems. They play a vital role in detecting signs of intrusion.

-
1. Assets generally include information, hardware, software, and people. Asset values are determined based on the impact to the organization if the asset is lost. Critical assets are those that are essential to meeting an organization's mission and business objectives. [Alberts 00] For this module, assets include information, hardware, and software that reside on and comprise the information technology infrastructure of an organization.
 2. Threat is defined here as anything that may compromise an asset. This could be a person, such as an employee or a hacker, or it could be a competitor or anyone else with deliberate intention to compromise an asset. Threats also include anything which results in accidental disruption to an asset (such as a natural disaster), the means of access to do so, or any outcome or consequence that results in an unwanted effect such as disclosure, modification, destruction, loss, or interruption. [Alberts 00]
 3. Systematic methods of information security risk analysis and assessment are emerging. These methods help an organization identify important assets, threats against these assets, security requirements for these assets, and weaknesses or vulnerabilities in current practice that increase the likelihood of these assets being compromised. Refer to *Operationally Critical Threat, Assets, and Vulnerability EvaluationSM(OCTAVESM) Framework, Version 1.0* [Alberts 99], *Secure Computing* [Summers 97], *Network Intrusion Detection: An Analyst's Handbook* [Northcutt 99], and "Web of Worries" [Kessler 00] for more information on this subject.
 4. Monitoring is the observation of data streams for specific events. Inspection is the examination of a data resource or process. Auditing is the systematic examination of data against documented expectations of form or behavior. Refer to *An Approach for Selecting and Specifying Tools for Information Survivability* [Firth 97].

This practice describes the topics your policy and procedures should address. They need to be tailored to reflect the specific business objectives and security requirements of your organization and its computing environment.

Why this is important

Having policy language and procedures in place that prepare you to detect signs of intrusion lets you use your procedures in a timely, managed, and controlled way and eliminates potential errors or omissions in advance of an attack. You do not want to be caught trying to determine what actions to take, what data to gather and preserve, and how to protect your systems from further damage while under attack or after the fact.

With advance planning, documentation, and education, trained staff members are able to more effectively coordinate their activities when detecting suspicious activity, an intrusion, or an intrusion attempt. Without the proper knowledge and skills, users may inadvertently expose parts of the organization to security threats.

How to do it

- *Include language in your organization's networked systems security policy that prepares you to detect signs of intrusion.*

Document the important and critical information assets and the level of protection (confidentiality, availability, integrity) required for each. Designations for the level may range from "cannot be compromised under any circumstances" (maximum protection) to "contains no sensitive information and can be easily restored" (minimal protection).

Document the types of threats or events that indicate possible signs of intrusion and also document how you intend to respond to them if they are detected. Types of threats may include⁵

- attempts (either failed or successful) to gain unauthorized access to a system or its data
- unintended and unauthorized disclosure of information
- unwanted disruption or denial of service
- the unauthorized use of a system to process, store, or transmit data
- changes to system hardware, firmware, or software characteristics without the knowledge or consent of you or of the asset owner

Recognize that there are threats that are difficult to protect against if your systems are connected to the Internet. You need to determine what actions you will take if these occur. Threats of this type include

- denial of service, including email bombing (sending a large volume of electronic messages to a targeted recipient until the system fails) and flood attacks (e.g., filling a channel with garbage, thereby denying others the ability to communicate across that channel or to the receiving host)
- programmed threats, such as new viruses not yet detected and eliminated by virus checking tools, and malicious software in the form of CGI scripts, plug-ins, servlets, or applets
- intruders probing or scanning your systems with the intent to exploit any vulnerabilities they discover for use in attempting an intrusion

5. Refer also to CERT/CC summaries, advisories, incident notes, and vulnerability notes available at <http://www.cert.org> and refer to "How To Eliminate The Ten Most Critical Internet Security Threats: The Experts Consensus" [SANS 00].

Document the requirement to establish and maintain secure, reliable configuration information for all assets that represent your known, expected state. This includes taking inventory and tagging all physical computing resources. Periodically compare this information with your current state to determine if anything has been altered in an unexpected way.

Document the roles, responsibilities, and authority of system administrators, security personnel, and users regarding the use and administration of all assets when they participate in detecting signs of suspicious behavior, including intrusions.

Document the roles, responsibilities, authority, and conditions for the testing of intrusion detection tools, the execution of intrusion detection procedures, and the examination of assets suspected of having been compromised. We strongly recommend that your policy requires that all such activity be conducted in a test environment isolated from production systems and networks.

➤ *Document procedures and take actions that implement your intrusion detection policy.*

In general terms, document what data you plan to collect⁶, why you want to collect it, and where and when you will collect it.

- Document what you want to discover by collecting the data. Typically, you should verify that levels of performance (function, throughput, load) are as you expected them to be and that there is an absence of errors and suspicious or unexpected behavior (as described in the subsequent practices of this module).
- Document where best to collect each type of data. We recommend collecting data as close to the source of its generation as possible, ideally on all hosts. If this is not possible because it would impact performance, collect the data as close to the host as possible. For example, if you cannot place an intrusion detection system on the firewall host, place it on another host that monitors network traffic on both the external (Internet) side of the firewall host and on the internal (organizational network) side of the firewall host.
- Document when to collect the data. As a starting point, we recommend collecting everything you possibly can, at every available source location, 24 hours a day, seven days a week. While this is likely to produce excessive amounts of data, you can manage this process by automating the deletion of data that you normally do not need to process or analyze further under normal operations. However, we recommend inserting a delay between the time of collection and the time of deletion⁷ in the event a suspicious event occurs that you want to analyze further using data you would normally delete.⁸

Document any special handling procedures for each type of collected data. This is particularly important for data that may be used as evidence in subsequent legal proceedings.⁹

-
6. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior.”
 7. A reasonable starting point might be one to two weeks, but this depends on your operation, review schedule, and data storage capacity.
 8. Refer to the practice “Manage logging and other data collection mechanisms,” specifically the step “Document your management plan for handling log files.”
 9. Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice “Collect and protect information associated with an intrusion.”

Document how you plan to conduct your review of all collected data. Because there is a large volume of system and network data that can be collected, and because there are increasing demands on an administrator's time, you need to carefully determine

- the order in which specific data will be reviewed
- the frequency of data review
- tools and other mechanisms, such as alerts, that can aid in better identifying suspicious and unexpected behaviors
- the types of events that warrant further investigation and in-depth analysis
- administrator authority, actions to be taken, and what circumstances warrant what actions
- how you will track the status of open events to resolution and closure

In particular

- document the procedure(s) by which monitoring is performed, i.e., the observation of data streams for specific events. This procedure specifies
 - the data streams to be monitored
 - the monitoring locations on systems and networks
 - the times and frequencies with which monitoring is to be performed
 - the activation of monitoring after the occurrence of what types of events
 - the operational activities necessary to alert appropriate personnel to act upon the suspected intrusion
- document the procedure(s) by which regular inspection and auditing of recorded data (e.g., logs) are performed to identify evidence of intrusions or intrusion attempts
- document the procedure by which physical audits of installed hardware and software are performed
- document the procedure by which integrity checking is performed (comparing the current operational state with a previously generated, secure, reliable, and known state). Specify
 - what files are to be checked
 - how integrity information is securely generated, maintained, and tested
 - frequency with which integrity checking is performed
- document the procedure by which correlation of intrusions is performed, i.e., determining when suspicious activity occurring in one part of your infrastructure may be related to activity in another part. Doing some level of correlation analysis during the intrusion detection process will assist you in determining the full extent of any compromise and its characteristics.¹⁰

Document the procedure for the acquisition and secure installation, configuration, and maintenance of all tools¹¹ necessary to implement your monitoring, inspection, auditing, and integrity checking procedures.

-
10. Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice "Analyze all available information to characterize an intrusion."
 11. One list of such tools is contained in the implementation "Identifying tools that aid in detecting signs of intrusion," available at <http://www.cert.org/security-improvement/implementations/i042.07.html>. Many of these tools can be downloaded from the Center for Education, Research, and Information Assurance Security [CERIAS] (formerly known as Computer Operations, Audit, and Security [COAST]), at <http://www.cerias.purdue.edu/>.

For each procedure and procedure step, document the roles, responsibilities, and authority of system administrators, security personnel, and users. Identify who performs each procedure activity, when, and under what conditions.

► *Conduct a legal review of your policy and procedures.*

This should be performed by your organization's legal counsel. It is intended to ensure that your policy and procedures

- are legally defensible and enforceable
- comply with overall company policies and procedures
- reflect known industry best practices demonstrating the exercise of due care
- conform to federal, state, and local laws and regulations
- protect your organization from being held legally responsible in the event of compromise
- require the preservation of critical evidence including a defensible, documented chain of custody for all artifacts that may be used in legal proceedings.¹²

► *Train the users who have authorized accounts on your systems.*

During the training process, users should learn

- what is expected of them
- how to identify suspicious behavior and who to notify
- what behaviors can reduce the exposure of systems to possible compromise, such as¹³
 - not opening unsolicited email attachments without verifying their source or checking their content
 - working with system administrators to install security patches for commonly used applications (such as web browsers)
 - not downloading and installing software from untrusted sources
 - making and testing backups
 - not using modems while connected through a local area network
 - protecting passwords and sensitive data
 - knowing how to respond to social engineering attempts
- what types of information are being gathered as part of routine security procedures, and the degree to which this information gathering may affect them.

Create and conduct periodic training on your intrusion preparation and detection policy and procedures. This training should be mandatory for all new employees and should cover aspects that are relevant to the employee's knowledge and responsibilities.

Test the effectiveness of the training and each employee's readiness. Conduct practice drills (e.g., detecting break-ins and viruses) that test procedures and execute operational activities, making sure all staff members are aware of their roles and responsibilities. Conduct post-mortem meetings with trainees. Provide remedial training as required.

12. Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice "Collect and protect information associated with an intrusion."

13. Refer to <http://www.sans.org> for more information on the five worst security mistakes committed by the average user.

Regularly conduct mandatory security awareness refresher training. Highlight recent changes to policy or procedures and summarize recent attack methods and counter measures. Make this subject a recurring topic at executive and management level staff meetings to maintain awareness.

To keep pace with the rapid rate of technological change, ensure that system and network administration staff have time set aside to maintain their knowledge, skills, and currency in technical topics required to implement your policy and procedures.

➤ *Keep your intrusion detection policy and all related procedures and training current.*

Periodically review your policy, procedures, and training. Take into account

- system changes and upgrades including the introduction of new software
- changes in critical assets
- changes in security requirements
- changes in key roles and responsibilities
- public and vendor information sources. These sources regularly report current intruder trends, new attack scenarios, security vulnerabilities, methods for their detection, and guidance to address them.

If your organization suffers an intrusion, review your policy, procedures, and training to determine if revisions are necessary to ensure that future intrusion attempts of the same type can be more readily detected and controlled, if not prevented.¹⁴

Other information

The most common sources of current information about security problems are the web sites of vendors, computer security organizations, and network security organizations. For example, you can find many advisories, incident notes, vulnerability notes, and tech tips at the CERT/CC web site (<http://www.cert.org>). Refer to the implementation “Maintaining currency by periodically reviewing public and vendor information sources,” available at <http://www.cert.org/security-improvement.implementations/i040.01.html>.

There are also mailing lists (such as those maintained by the SANS Institute with subscriptions available at <http://www.sans.org>), some of which are sponsored by vendors, and USENET news groups. Because lists and web sites appear, disappear, change frequently, or cease to be updated regularly, you need to ensure that the sources you consult are up-to-date.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL

<http://www.cert.org/security-improvement/practices/p090.html>

14. Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice “Identify and implement security lessons learned.”

2

Identify data that characterize systems and aid in detecting signs of suspicious behavior.¹

Collecting data generated by system, network, application, and user² activities is essential for analyzing the security of your information assets and detecting signs of suspicious and unexpected behavior. Log files contain information about past activities. You should identify the logging mechanisms and types of logs (system, file access, process, network, application-specific, etc.) available for each asset and identify the data recorded within each log.

Different systems provide various types of logging information; some systems do not collect adequate information in their default condition. It is important to supplement your logs with additional collection mechanisms that watch for signs of intrusions or intrusion attempts. They should also alert responsible parties when events occur. Include mechanisms that

- monitor and inspect system resource use
- monitor and inspect network traffic and connections
- monitor and inspect user account and file access
- scan for viruses
- verify file and data integrity
- probe for system and network vulnerabilities
- reduce, scan, monitor, and inspect your log files

Capturing an accurate, reliable, and complete characterization of your systems when they are first created, and as they evolve, establishes the expected state against which to compare your current systems. The information to be captured includes a known, expected state for all assets, including your network traffic, system and network performance, processes, users, files and directories, and hardware. This includes information that characterizes past behavior derived from system logs and monitoring tools, which is available once you have been operational for some period of time. This trusted record is periodically compared with your current systems to determine if assets are behaving as expected; in other words, to verify the integrity of your systems and to identify any deviations from expected behavior.

Characterizing your software, hardware, and information assets is a time-consuming, complex, and ongoing task. You need to determine, in advance, the level of resources you can commit to this activity.

-
1. includes signs of intrusion and intrusion attempts
 2. Users are those who access, administer, and manage your systems and have authorized accounts on your systems.

Why this is important	<p>Approaches to detecting signs of suspicious or unexpected behavior are often based on identifying differences between your current operational state and a previously captured and trusted expected state.</p> <p>You need to know where each asset is located and what information you expect to find in each location. You need to be able to verify the correct or expected state of every asset. Without this information, you cannot adequately determine if anything has been added, deleted, modified, lost, or stolen.</p> <p>You may not be able to rebuild a critical component that has been compromised without up-to-date, available, trusted characterizations.</p> <p>Log files may be the only record of suspicious behavior. Failure to enable the mechanisms to record this information and use them to initiate alert mechanisms will greatly weaken or eliminate your ability to detect intrusion attempts and to determine whether or not they succeeded. Similar problems can result from not having the necessary procedures and mechanisms in place to process and analyze your log files.</p> <p>You may need your logs to</p> <ul style="list-style-type: none">• alert you to suspicious activity that requires further investigation• determine the extent of an intruder's activity• help you recover your systems• provide information required for legal proceedings <p>It is possible that the logging and monitoring mechanisms provided with your systems may not produce all of the information necessary to detect signs of an intrusion in a timely manner. If adequate information is provided, the volume of data may be so overwhelming that automated analysis to reduce it to a manageable subset is required before you can examine it for signs of intrusive activity. In either case, you will need to add tools to your systems to adequately detect signs of suspicious or unexpected behavior that require further analysis.³</p>
------------------------------	--

How to do it	<p>► <i>Determine what data is most useful to collect.</i></p> <p>You need to balance the importance of recording system, network, and user activities with the resources available to store, process, review, and secure them. Questions that help you determine the usefulness of collected data include</p> <ul style="list-style-type: none">• What is the priority of this asset (hardware, software, information)? How important is it to collect data related to this asset? How important is it to characterize this asset?• What is the system's sole or primary purpose? For example, if a host is acting as a web server, you want to capture web logs.• How many users are assigned to the system and how important is it for you to know who is logged on? This helps you decide how much login/logout information to capture.
---------------------	---

3. Refer to the practices "Monitor and inspect network activities for unexpected behavior," "Monitor and inspect system activities for unexpected behavior," and "Inspect files and directories for unexpected changes" for examples of suspicious and unexpected behaviors that can be determined based on the data you collect.

- How important is it to be able to use your logs and other data to recover a compromised system? This helps you set the priority for capturing information such as data and file transaction logs.
- What are the range of services that can be performed on this system? Process accounting information is useful to detect unauthorized services and intruder actions.
- What is your organization's ability and capacity to process and analyze all collected data to obtain useful information when it is needed?

► *Identify the data to be collected.*

A table of data categories and possible types of data to collect is shown below.

Table 1: Data Categories and Types of Data to Collect

Data Category	Types of data to collect
Network performance	<ul style="list-style-type: none"> • total traffic load in and out over time (packet, byte, and connection counts) and by event (such as new product or service release) • traffic load (percentage of packets, bytes, connections) in and out over time sorted by protocol, source address, destination address, other packet header data • error counts on all network interfaces
Other network data	<ul style="list-style-type: none"> • service initiation requests • name of the user/host requesting the service • network traffic (packet headers) • successful connections and connection attempts (protocol, port, source, destination, time) • connection duration • connection flow (sequence of packets from initiation to termination) • states associated with network interfaces (up, down) • network sockets currently open • whether or not network interface card is in promiscuous mode • network probes and scans • results of administrator probes
System performance	<ul style="list-style-type: none"> • total resource use over time (CPU, memory [used, free], disk [used, free]) • status and errors reported by systems and hardware devices • changes in system status, including shutdowns and restarts • file system status (where mounted, free space by partition, open files, biggest file) over time and at specific times • file system warnings (low freespace, too many open files, file exceeding allocated size) • disk counters (input/output, queue lengths) over time and at specific times • hardware availability (modems, network interface cards, memory)

Data Category	Types of data to collect
Other system data	<ul style="list-style-type: none"> • actions requiring special privileges • successful and failed logins • modem activities • presence of new services and devices • configuration of resources and devices
Process performance	<ul style="list-style-type: none"> • amount of resources used (CPU, memory, disk, time) by specific processes over time; top “x” resource-consuming processes • system and user processes and services executing at any given time
Other process data	<ul style="list-style-type: none"> • user executing the process • process start-up time, arguments, file names • process exit status, time, duration, resources consumed • the means by which each process is normally initiated (administrator, other users, other programs or processes), with what authorization and privileges • devices used by specific processes • files currently open by specific processes
Files and directories	<ul style="list-style-type: none"> • list of files, directories, attributes • cryptographic checksums for all files and directories • accesses (open, create, modify, execute, delete), time, date • changes to sizes, contents, protections, types, locations • changes to access control lists on system tools • additions and deletions of files and directories • results of virus scanners
Users	<ul style="list-style-type: none"> • login/logout information (location, time): successful attempts, failed attempts, attempted logins to privileged accounts • login/logout information on remote access servers that appears in modem logs • changes in user identity • changes in authentication status, such as enabling privileges • failed attempts to access restricted information (such as password files) • keystroke monitoring logs • violations of user quotas

Data Category	Types of data to collect
Applications	<ul style="list-style-type: none"> applications- and services-specific information such as network traffic (packet content), mail logs, FTP logs, web server logs, modem logs, firewall logs, SNMP logs, DNS logs, intrusion detection system logs, database management system logs. Services specific information could be <ul style="list-style-type: none"> for FTP requests: files transferred and connection statistics for web requests: pages accessed, credentials of the requestor, connection statistics, user requests over time, which pages are most requested, and who is requesting them for mail requests: sender, receiver, size, and tracing information; for a mail server, number of messages over time, number of queued messages for DNS requests: questions, answers, and zone transfers for a file system server: file transfers over time for a database server: transactions over time
Log files	<ul style="list-style-type: none"> results of scanning, filtering, and reducing log file contents checks for log file consistency (increasing file size over time, use of consecutive, increasing time stamps with no gaps)
Vulnerabilities	<ul style="list-style-type: none"> results of vulnerability scanners (presence of known vulnerabilities) vulnerability patch logging

➤ *Identify the data to be captured using logging mechanisms.*

Identify the

- types of information you can log
- mechanisms used for logging
- locations where the logging is performed
- locations where the log files are stored

Use Table 1 as a guide to the types of information to log (although not all systems are able to log every type in the table). Tailor logging selections to meet your site's specific policies and security requirements.

For all data categories, capture alerts and any reported errors.

If possible, do not log passwords, even incorrect ones. Logging correct passwords creates an enormous potential vulnerability if a non-authorized user or intruder accesses log files. Recording incorrect passwords is also risky as they often differ from valid passwords by only a single character or transposition. Turning off password logging may require resetting a system default. If you cannot turn off password logging, you need to exercise special care in protecting access to log files that contain this information.⁴ However, you may want to log data about password use, such as the number of failed attempts, accesses to specific accounts, etc.

4. Refer to the practice "Manage logging and other data collection mechanisms."

- *Determine if the logging mechanisms provided with your systems sufficiently capture the required information.*

Determine the logging mechanisms available for the systems at your site. Identify what types of information each logging mechanism can capture. There may be differences in the log file contents provided by different vendors, even for similar types of systems.

Determine where each logging mechanism stores data. Identify how the log files are named and where they are located. The names of these log files can differ even among versions of the same operating system delivered by a single vendor, so it is important that you verify this each time you upgrade your systems.

- *Identify the data to be captured using additional data collection mechanisms (such as monitoring).*

Use Table 1 as a guide to the types of information you need to collect beyond that which is available using logging mechanisms. Tailor additional data selections to meet your site's intrusion detection policies and procedures.

Monitoring is the observation of data streams for specific events, whereas logging systematically records specified events in the order that they occur. Monitoring generally connotes more of a "real time" analysis activity, while inspecting⁵ log files generally occurs as more of an "offline" or after-the-fact activity. Monitoring is often preferable where there are large quantities of data, such as network traffic. In most circumstances, it isn't feasible to store every network packet but monitoring the network traffic for specific types of events and connections is very desirable.

Real-time intrusion detection systems⁶, including log file monitoring tools, can detect possible intrusions or access violations as they are occurring and generate alerts in any of the Table 1 data categories. Real-time intrusion detection occurs while an intruder is attempting to break in or is still present on your system. This is contrasted with offline intrusion detection which is performed after the intrusion has occurred, usually through inspecting various system and network log files and performing data and system integrity tests.

A host- or system-based intrusion detection system (IDS) examines data such as log files, process accounting information, and user behavior and generates alerts based on specified configuration information. A network IDS examines network traffic, including packet headers and content.

Both types of ID systems can employ one or more analysis approaches to determine whether or not an intrusion has occurred. The two most common analysis approaches are

- attack signature detection (sometime called "misuse detection") which identifies patterns (signatures) corresponding to known attacks
- anomaly detection, which identifies any unacceptable deviation from expected behavior. Expected behavior is defined, in advance, by a manually or automatically developed profile of system, network, and user behavior.

It is difficult to provide guidance about additional data selection and collection mechanisms because the selection criteria varies based on organizational policy and security requirements. This is made more complex by a lack of uniformity in the intrusion characterizations used by common collection mechanisms.

5. Inspection is the examination of a data resource or process.

6. Refer to *State of the Practice in Intrusion Detection Technologies* [Allen 00].

In most cases, you will need to perform manual analysis in concert with the automated data collection and reporting performed by any mechanism.

➤ *Determine which events should produce an alert.*

Events that require immediate administrator attention and that need to be given the highest priority should be designated as alerts. Alerts can be in the form of a message displayed on your workstation, someone stopping by your office, a phone call or voice mail, email, or pager messages. Most data collection mechanisms provide some form of alerting capability for specified events.

➤ *Recognize the iterative nature of data collection and characterization.*

Initially, you want to collect as much data as possible because you do not know which data will be the most meaningful. Over time, you begin to identify filtering approaches that allow you to successively refine what data you choose to examine. In addition, you begin to identify trends in behavior and specific activities and events that constitute normal behavior. This behavior can then be captured as part of your current characterization baseline (see the following steps). However, it is important to note that systems and the operations that users perform on them are always changing (for example, the addition of new software, or changes in user privileges due to new assignments), so you need to periodically examine your characterization information to determine if it needs to be updated.

Having a trusted characterization baseline is critical for identifying departures from normal and expected behavior that warrant further investigation. In addition, data that reveals normal and expected behavior can be eliminated from further consideration, allowing an administrator to focus attention on a smaller set of data demonstrating any unexpected behavior, including potential intrusions.

➤ *Document and verify your characterization trust assumptions.*

As you generate all characterization information (for both baselining and comparison purposes), explicitly document your trust assumptions and continually verify that you can trust the results produced.

Trust assumptions will likely address

- the operating system kernel (loaded from virus-free, secure distribution media)
- the media where your data collection and characterization tools are stored and from which they are installed
- cryptographic checksums and authoritative reference data that constitute characterization data

➤ *Characterize typical network traffic and performance.*

Document the procedure by which you intend to verify that the network traffic traversing your networks is as expected and reflects, for example, trusted source and destination addresses, and legitimate ports and protocols.

The types of network traffic information that you should capture includes network performance and other network data described in Table 1 and answers the following questions:

- What traffic is typically produced by my system?

- What traffic is typically consumed by my system?
- What are the range of acceptable performance levels provided by my networks?

Comparing previous network performance information with current information allows you to determine if any network performance characteristic is beyond tolerable or acceptable limits.⁷

➤ *Characterize expected system behavior and performance.*

Document the procedure by which you intend to verify that your systems are performing as expected.

The types of performance information you want to capture answers the question “What is the range of acceptable performance levels provided by my systems?” and includes system performance data and other system data described in Table 1.

Comparing previous system performance information with current information allows you to determine if any system performance characteristic is beyond tolerable or acceptable limits.⁸

➤ *Characterize expected process and user behavior.⁹*

Document the procedure by which you intend to verify that the processes executing on your systems are operating only as expected and attributed only to authorized activities of users, administrators, and system functions.

The types of process information you want to capture answers the question “What processes are normally running on my system?” and includes process and user data described in Table 1.

Comparing previous process and user information with current information allows you to determine if any process is behaving in an unexpected or suspicious manner.

➤ *Characterize expected file and directory information.*

Document the procedure you will use to verify that the files and directories on your systems are as you expect them to be and that they were created, modified, accessed and deleted as you expected.

For each file and directory, the type of information you want to capture should include file and directory data described in Table 1 and answer the following questions:

- What files are on my system (name, type, attributes, etc.) and where do they reside?
- How are files and directories affected during normal system operation (created, deleted, contents changed, accessed, permissions changed, location changed)?

7. Refer to the practice “Monitor and inspect network activities for unexpected behavior.”
8. Refer to the practice “Monitor and inspect system activities for unexpected behavior.”
9. Refer to the practice “Monitor and inspect system activities for unexpected behavior.” Also refer to the implementation “Process analysis checklist” available at <http://www.cert.org/security-improvement/implementations/i005.02.html>.

Capture a cryptographic checksum for all files and directories. For example, *Tripwire*¹⁰ will generate this as well as inform you of the state of the collection of files on your system (added or deleted), changes in state (protection changes), and the fact that changes to file contents have or have not occurred (but not what the actual changes are).

Commercial versions of *Tripwire* are available for UNIX and Windows NT systems. MD5¹¹ and other one-way hashing functions (such as SHA-1, RIPEMD-160, and HAVAL¹²) can also be used to generate cryptographic checksums.

Important files and directories to characterize include

- operating systems and configuration files
- access control lists
- applications
- security tools and data such as those used for integrity checking and detecting signs of intrusion
- organizational data such as financial reports and employee information
- user data
- public information such as web pages

Some operating systems provide the capability to make files *immutable*, meaning unchangeable by *any* process on the system, including system and administrative processes. All operating system and other files that don't need to be modified when a system is running should be made immutable, where possible.

Comparing previous file and directory information with current information allows you to determine if any file or directory has changed in an unexpected or suspicious manner.¹³

► *Generate an inventory of your system hardware.*

If you have not already done so, create an inventory of all of your computing hardware assets. This is most likely accomplished by performing a physical audit. Use a tool (e.g., a database management system or spreadsheet) to record the initial inventory and keep it up-to-date. Select a tool that will easily allow you to perform comparisons with subsequent inventories.

Ensure that procedures are in place to update your hardware inventory when the physical location of equipment changes, when its hardware configuration is upgraded (e.g., memory is added), and when equipment is added to or removed from your systems.

Produce and maintain complete, up-to-date network infrastructure information that captures the architecture, connectivity, and identity of all network devices.

-
10. Refer to <http://www.tripwiresecurity.com> and the implementation “Installing, configuring, and using Tripwire to verify the integrity of directories and files on systems running Solaris 2.x.” available at <http://www.cert.org/security-improvement/implementations/i002.02.html>.
 11. Refer to the implementation “Using MD5 to verify the integrity of file contents” available at <http://www.cert.org/security-improvement/implementations/i002.01.html>.
 12. References for SHA-1 (Secure Hashing Algorithm), RIPEMD-160, and HAVAL (Hashing Algorithm with VAriable Length) can be found at <http://www.users.zetnet.co.uk/hopwood/crypto/scan/md.html>.
 13. Refer to the practice “Inspect files and directories for unexpected changes.”

This includes¹⁴

- the layout or topology of all network devices
- network architecture
- network and device connectivity
- network and device configuration
- administrative domains
- physical location of all network devices
- intermediate public networks, if any

Identify network monitoring and management mechanisms to keep this information up-to-date and to alert you to anomalies.

Use automated tools to detect installed hardware and compare the results with your physical inventory. For PC-based systems, the Windows95, 98, or NT operating systems provide a complete hardware inventory capability as part of system properties. There are also a variety of vendor and public domain tools available such as *nmap*¹⁵. Tools such as *daemon dialers* can help determine what modems are connected to your telephone lines, systems, and networks.

Refer to the implementation “Establishing and maintaining a physical inventory of your computing equipment,” available at <http://www.cert.org/security-improvement/implementations/i043.02.html>.

► *Protect your asset characterization information, authoritative reference data, and hardware inventory to ensure their integrity.*

Keep authoritative reference copies of files and checksums on write-protected or read-only media stored in a physically secure location. You may want to consider using a tool such as PGP (Pretty Good Privacy) to “sign” the output generated by your checksum tool.

Consider making paper copies of configuration files and cryptographic checksums in the event you are unable to recover uncorrupted electronic versions.

If you transmit authoritative reference data over unsecured network connections, make sure to verify the data upon arrival at the destination host (e.g., by using MD5). Consider encrypting the reference data at the source host to reduce the likelihood of the information being compromised, to protect confidentiality and privacy, and to prevent password capture.

Encrypt your asset characterization information, authoritative reference data, and hardware inventory if your organization’s security requirements demand this level of protection.

► *Keep your asset characterization information, authoritative reference data, and system inventory up to date.*

14. This includes details such as MAC addresses, IP addresses, host names, ports on routers, hubs and switches, contact telephone numbers, and external devices and servers your network depends on such as ISP routes, DNS servers, and WWW cache.

15. Available at <http://www.insecure.org>.

Policy considerations	Your organization's networked systems security policy should require that your system administrators create an accurate, reliable, and complete characterization of those assets you have selected when they are first created and at well-defined events when you modify, add to, and replace elements of your systems or determine that the characterization of normal, expected behavior needs to change.																
Other information	<ol style="list-style-type: none"> 1. It is difficult to estimate both the time required to develop an initial characterization baseline and the additional time required to keep it updated. One good guideline is to periodically observe a host's behavior for three to six months and then derive the initial characterization baseline from that observation, using some of the data collection mechanisms described in this practice. Another good guideline is to allow 15 days of observation for characterizing the first host, 12 days for the second host, 9 days for the third host, and 7 days for the fourth and all subsequent hosts. It may take one year or more for an administrator to observe the network traffic traversing a large network and to develop a characterization baseline representing normal traffic behavior. Once an administrator understands how to develop a characterization baseline, developing subsequent baselines should proceed more quickly. 2. For information on log filtering, analysis, and alerting approaches, see the module <i>Responding to Intrusions</i> [Kossakowski 99] and the supporting implementations for <i>Detecting Signs of Intrusion</i> [Allen 00]. These are available at http://www.cert.org/security-improvement/implementations. See also <i>State of the Practice of Intrusion Detection Technologies</i> [Allen 99]. 3. Tool selection:¹⁶ <p>You may find it useful to categorize and select tools using a set of activities associated with common approaches for detecting signs of suspicious or unexpected behavior. Such a set of activities would include</p> <table> <tbody> <tr> <td>filtering</td> <td>examining a data stream and removing from it items that are deemed undesirable or inappropriate</td> </tr> <tr> <td>probing</td> <td>attempting connections or queries</td> </tr> <tr> <td>scanning</td> <td>iteratively probing a collection of systems or data for known vulnerabilities</td> </tr> <tr> <td>monitoring</td> <td>observing a data stream for specified events</td> </tr> <tr> <td>inspecting</td> <td>examining a data resource or process</td> </tr> <tr> <td>auditing</td> <td>systematically examining system data against documented expectations of form or behavior</td> </tr> <tr> <td>integrity checking</td> <td>verifying that the contents of a data resource are exactly as created, stored, or transmitted</td> </tr> <tr> <td>notifying</td> <td>alerting a designated recipient to the occurrence of a specific event</td> </tr> </tbody> </table>	filtering	examining a data stream and removing from it items that are deemed undesirable or inappropriate	probing	attempting connections or queries	scanning	iteratively probing a collection of systems or data for known vulnerabilities	monitoring	observing a data stream for specified events	inspecting	examining a data resource or process	auditing	systematically examining system data against documented expectations of form or behavior	integrity checking	verifying that the contents of a data resource are exactly as created, stored, or transmitted	notifying	alerting a designated recipient to the occurrence of a specific event
filtering	examining a data stream and removing from it items that are deemed undesirable or inappropriate																
probing	attempting connections or queries																
scanning	iteratively probing a collection of systems or data for known vulnerabilities																
monitoring	observing a data stream for specified events																
inspecting	examining a data resource or process																
auditing	systematically examining system data against documented expectations of form or behavior																
integrity checking	verifying that the contents of a data resource are exactly as created, stored, or transmitted																
notifying	alerting a designated recipient to the occurrence of a specific event																

16. Refer to *An Approach for Selecting and Specifying Tools for Information Survivability* [Firth 97].

4. One list of monitoring and intrusion detection tools is contained in the implementation “Identifying tools that aid in detecting signs of intrusion,” available at <http://www.cert.org/security-improvement/implementations/i042.07.html>. Many of these tools can be downloaded from the Center for Education, Research, and Information Assurance Security [CERIAS] (formerly known as Computer Operations, Audit, and Security [COAST]), at <http://www.cerias.purdue.edu/>.
5. Some guidance on evaluation criteria for selecting monitoring systems and tools is available in “System and Network Monitoring” [Sellens 00]. Criteria include
 - size and complexity
 - scalability
 - reliability
 - cost
 - number and type of probes
 - configuration complexity and flexibility
 - exception reporting style
 - exception reporting tools
 - logging and data storage
 - reporting mechanisms

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL

<http://www.cert.org/security-improvement/practices/p091.html>

3

Manage logging and other data collection mechanisms.

Once you have identified¹ the data to be collected, you need to enable the corresponding logging and data collection mechanisms as well as log filters and alert tools. All of these mechanisms can produce a large volume of recorded information. You need to determine how to best capture, manage, and protect all recorded information, and determine how to alert security staff and administrators when appropriate.

Select and enable data collection mechanisms based on your site's security policy and security requirements.

Why this is important

Failure to enable the necessary data collection mechanisms will greatly weaken or eliminate your ability to detect suspicious behavior and intrusion attempts and to determine whether or not such attempts succeeded.

Failure to configure and secure the volume of data produced by these mechanisms will place the data at risk of compromise and make subsequent review and analysis difficult, if not impossible.

How to do it

➤ *Enable logging.*

Using the logging mechanisms provided by the vendor and any supplemental tools, enable all logging that you have selected. For help, refer to the administration documentation for your systems to learn how to enable each of the logging mechanisms and refer to any documentation that accompanies relevant tools. This documentation will specify whether these mechanisms need to be enabled only once, each time the system is rebooted, or at regular intervals during the system's normal operation. Some logging mechanisms let you select different levels of detail.

Pay attention to the location of the log data: some tools allow you to choose a file or directory where the data is logged while others write their data to a predefined, default location. Make sure that you have sufficient space for the data that is generated. Ensure that the logged data is protected, based on previously determined ACLs (access control lists) and your security policy.

1. Refer to the practice "Identify data that characterize systems and aid in detecting signs of suspicious behavior."

Be aware that multiple logging mechanisms may contribute log records to a single log file, such as syslog in UNIX systems. This is specified within your system configuration file.²

► *Protect logs to ensure they are reliable.*

To protect sensitive information, ensure that log files are protected from being accessed or modified by unauthorized users. Confirm that only authorized users can access utilities that reconfigure logging mechanisms, turn the utilities on and off, and write to, modify, and read log data.

It is important to collect and archive log files so that they cannot be accessed by an intruder to remove or alter signs of an intrusion or add erroneous information. Use the following methods to ensure log files are not modified:

- Log data to a file on a separate host that is dedicated solely to log collecting. The log host should reside in a physically secure location that is not easily accessible from the network. For example, capturing log data using a computer via a dedicated serial line provides a way of storing the log files more securely than if they were written on the logging host's disks.
- Log selected data to a write-once/read-many device (such as CD-ROM or a specially configured tape drive) or to a write-only device (such as a printer) to eliminate the possibility of the data being modified once it is written.
- If supported by your systems, set selected log file attributes that enable only new information to be appended to the log files (i.e., new records can be added, those already recorded cannot be modified).
- Encrypt log files, particularly those that contain sensitive data or those being transmitted across a network.

Logging directly to disk on the local host is easiest to configure and allows instant access to file records for analysis, but it is also the least secure. Collecting log files on a write-once device requires slightly more effort to configure but is more secure. However, data is not as easily accessed and you need to maintain a supply of storage media.

Printing the logging results is useful when you require permanent and immediate log files, but printed logs can be difficult to search, require manual analysis, and require a potentially large storage space.

When the host generating the logging data is different from the host recording it, you must secure the path between them. For environments where short distances separate the generating host from the recording host, you can connect them with single point-to-point cable(s). For environments where this approach is not practical, minimize the number of networks and routers used to make the connection or encrypt sensitive log data as it is generated.

To protect the log files on your log host, place the host on a separate, secure subnet that is protected by a firewall and make log files "read only" from the log host console.

You need to prepare systems that perform logging to ensure that they do not stop functioning in the event of a logging denial-of-service attack. For UNIX systems, an intruder could launch an attack that fills up the syslog files so that when the logging partition is full, logging ceases. For NT systems, an intruder could overwrite the oldest log file entries after filling all available storage.

2. Refer to the implementation "Understanding system log files on a Solaris 2.x operating system" available at <http://www.cert.org/security-improvement/implementations/i041.12.html>.

To prepare a system so it continues functioning, create separate file partitions for different log information and filter network messages³ to decrease the likelihood of such attacks.

In addition, some systems provide the capability to shut down (or prohibit anyone but the system administrator to login), and produce a warning when the log files are full. However, this is not normally the default configuration so it must be explicitly specified.

► *Document your management plan for handling log files.*

Handle the total volume of logged information. We recommend that you log as much as possible for your systems and networks. While log files can very quickly consume a great deal of storage (which is relatively inexpensive), it is difficult to anticipate which logs will be critical in the event of an intrusion. Based on your log collection and storage approach, you may want to compress log files to allow them to remain accessible online for easier review and to conserve space.

Rotate log files. This activity consists of

- making a copy of the active (online) log files at regular intervals (ranging from daily to weekly)
- renaming the files so information contained in the file is not further augmented
- resetting file contents
- verifying that logging still works

Rotating log files allows you to limit the volume of log data you have to examine at any given time. It also allows you to keep log files open for a limited duration so that damage is bounded if an active log file is compromised. In this way, you create a collection of log files that contain well-defined time intervals of recorded data.

You can then consolidate logs from different systems by matching time intervals. This will help you gain a network-wide perspective on the activities. To perform this consolidation, you will likely need to merge log files from different systems into a central log file.⁴ To avoid having to adjust the timestamps used in each, use a master clock system such NTP (Network Time Protocol) or another time synchronization protocol system.⁵ Make sure to take into account different time zones and formats for recorded time.

Back up and archive log files. Move your log files to permanent storage or capture them as part of your regular backup procedure. This will allow you to retrieve them later if the need arises. Document the method you use to access archived log files. Create backups before you execute any automated tools that truncate and reset the log files so that minimal logging data is lost.

Encrypt log files. We recommend encrypting log files that contain sensitive data as the log data is being recorded. Protect the encryption software and place a copy of your encryption keys on a floppy disk or WORM (write once, read many) CD-ROM in a secure location such as a safe or safety deposit box. If the keys are lost, the log files cannot be used. If possible, use public key encryption.

3. Refer to the module *Deploying Firewalls* [Fithen 99], specifically the practice “Configure firewall packet filtering.”

4. Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice “Analyze all available information to characterize an intrusion” and the step “Examine logs generated by firewalls, network monitors, and routers.”

5. Refer to <http://www.eecis.udel.edu/~ntp/>.

The logs can be encrypted using the public key (which can be safely stored online) and the corresponding private key (stored offline) can then be used to decrypt the logs.

Ensure that you have the system and personnel resources necessary to analyze logs on a regular basis (at least daily in most cases) and on demand (such as when alert events occur).

Dispose of log files. Ensure that all media containing log file data are securely disposed (e.g., shredding hardcopy output, sanitizing disks, destroying CD-ROMs).

► *Protect data collection mechanisms and their outputs to ensure they are reliable.*

Make sure you obtain tools from a reliable source and verify the software integrity through digital signatures, cryptographic checksums, or by using trusted copies from secure media. Intruders have been known to modify tools installed by authorized administrators so that the tools, when used, do not identify the presence of the intruder.

Once you have verified the software, you need to configure it for use at your site. The installation should be performed on a secure system to eliminate the possibility of the tool being tampered with before you have had a chance to deploy it. You should make a cryptographic checksum of these tools. Using this information, you can then verify that your original configuration has not been compromised. You need to protect these tools by ensuring that they have the appropriate access control lists set to allow use and modification only by authorized users. The reports produced by these tools also need to be protected so only authorized users can use them.

► *Take into account special data collection and handling procedures required to preserve data as evidence.*

This is required in the event an intrusion has actually occurred and your organization decides to take legal action against the intruder.

Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice “Collect and protect information associated with an intrusion.”

Policy considerations

Your organization’s security policy for networked systems should

- require that you document a management plan for handling log files. This plan should include what to log, when and why to log, where to log, and who is responsible for all aspects of the plan.
- identify approved sources for acquiring tool software (Internet, shareware, purchased from vendor, etc.) and acceptable use practices related to tools

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL

<http://www.cert.org/security-improvement/practices/p092.html>

4

Ensure that the software used to examine systems has not been compromised.

When you look for signs of intrusions on your systems, and when you examine your systems in general, you should use a verified, reference set of software—one that contains only trusted copies of software that have not been modified—and perform a clean boot (start the system from a known, virus-free image of the operating system). In addition to executable programs, the verified set of software must include all the operating system kernel, system libraries, configuration and data files, and system utilities on which the programs depend. You should avoid relying on software that resides on systems being examined (unless you can verify that the software and its supporting libraries, configuration files, and data files have not been modified).

Why this is important

Intrusion detection depends heavily on the reliability of the information you gather about the state and behavior of your systems. Therefore, it is essential that you use only software that you know to be reliable and accurate in its reporting of such information.

Intruders often replace software that would reveal their presence with substitutes that obscure or remove such information. Intruders are known to have replaced programs, libraries, and other utilities called by the programs. If a program used in detecting intrusions has been tampered with or replaced with a substitute, obviously you cannot rely on its output.

Ensuring that you are using only verified software may be very difficult. Intruders can make extremely devious system modifications that make things appear to be normal when, in fact, they are not. For example, using the rootkit tool set¹, an intruder can replace the *ps* command on a UNIX system with one that does not display the intruder's process, or an editor can be replaced with one that reads a file other than the one specified (the intruder may have hidden the original file and replaced it with another version). Intruders may also modify software that is executed at system boot and shutdown, complicating your ability to safely take a system offline for more detailed analysis. Viruses often do this.

How to do it

Note: Any examination or alteration of a suspect system could destroy data that may be useful during any legal investigation or proceedings. However, to determine the cause of the problem and return a system to operations as soon as possible, the system administrator may have no choice but to destroy such data.

1. For Windows NT, refer to <http://www.rootkit.com>. For Linux, refer to <http://www.securityfocus.com/tools/1489>. To check for signs of the presence of rootkit, refer to <http://www.securityfocus.com/tools/1646>.

If you require that legal evidence be preserved, we recommend that you initiate your intrusion response procedures immediately, as described in *Responding to Intrusions* [Kossakowski 99].

The guiding principle for this practice is that you maintain a certain amount of suspicion. Question everything you observe, and answer the questions

- What software is producing this output?
- What other software does it rely on?
- What software can I trust?

There are five general alternative approaches to achieve the goal of using a verified set of software. Each approach has advantages and disadvantages so you should choose a method appropriate to your current circumstances.

In all cases, the verified software should be located on physically write-protected media (e.g., CD-ROM or write-protected disk), so that it cannot be modified by a user or by software running on the system being examined.

- *Move the disk from the system suspected of having been compromised to a write-protected, verified system and examine the disk's contents using the software on the verified system.*

The advantage of this method is that you do not need to rely on the integrity of any part of the operating system or the hardware on the suspect system. The method is effective and reasonable when you suspect that a particular system has been compromised and you want to analyze it. However, it may not be practical for automated procedures or for checking a large number of systems.

Be careful when shutting down the suspect system since the mere act of doing so may result in hiding or losing the evidence you are seeking. Before shutting down the suspect system, look at any programs that will run at shutdown for signs that they were modified (for example in some UNIX operating systems, the */etc/shutdown* program should be examined). However, be aware that just looking at the file may be misleading since you are relying on the software on the suspect system. You may want to execute verified copies of shutdown programs and their data files (taking care to save the original files for later analysis). Other alternatives are to execute the shutdown from external media, force the system to halt immediately, or to just pull the plug.

- *Attach to the suspect system a write-protected, verified system disk that contains the operating system and all needed software, and then reboot the system using the verified operating system.*

This method has similar advantages and disadvantages to those of the method above but relies on the trustworthiness of the hardware of the suspect system.

- *Generate an image of the suspect system disk, mount it on a verified system, and examine it there.*

This method is acceptable if you have a verified system that you can use to examine the suspect system disk.

Advantages of this approach include not affecting the operational environment of the suspect system because you're examining an image of it on another system and preserving the original evidence for subsequent legal proceedings.

► *Use external media containing a verified set of software to examine the suspect system.*

To use this method, you need a CD-ROM or a write-protected disk containing verified software to be used when examining the suspect system.

A significant concern with this approach is that you will still be using the operating system of the suspect system (e.g., the UNIX kernel), and it is highly unlikely that you have provided every needed program, utility, and library. As a result, the outcome of such analysis is suspect.

► *Verify the software on the suspect system first, then use the verified software to examine the suspect system.*

This method requires that you compare the software on the suspect system with a reference copy (either complete files or cryptographic checksums).² However, take care to use a verified comparison program or cryptographic checksumming program. The program used to verify the software should be located on physically write-protected media.

This approach has the same problem as that noted in the previous step with respect to using the operating system of the suspect system.

Policy considerations	Your organization's networked systems security policy should specify the level of verification that is required when examining each class of data and service provided by the organization's systems.
Other information	<ol style="list-style-type: none">1. Some operating systems provide the capability to make files <i>immutable</i>, meaning unchangeable by <i>any</i> process on the system including system and administrative processes. All operating system files that don't need to be modified when a system is running should be made immutable, where possible.2. When you are examining your system through a remote access connection, you need to be sure that you have established a secure channel to the system (such as that provided by ssh³ [secure shell]) so that only authorized personnel use the channel and nothing is changed or revealed in transit.
Where to find updates	The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL http://www.cert.org/security-improvement/practices/p093.html

2. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior.”

3. Refer to the module *Securing Network Servers* [Allen 00], specifically the practice “Configure computers for secure remote administration” and the supporting implementation “Installing, configuring, and operating the secure shell (ssh) on systems running Solaris 2.x,” available at <http://www.cert.org/security-improvement/implementations/i062.01.html>.

5

Monitor and inspect network activities for unexpected behavior.

Data about network activities (traffic, performance, etc.) can be collected from a variety of sources such as

- administrator probes (ICMP¹ pings, port probes, SNMP² queries)
- log files (routers, firewalls, other network hosts and devices)
- alert reports
- error reports
- network performance statistics reports
- the outputs of tools used to support in-depth analysis

Unexpected network behavior you should watch for includes

- unexpected changes in network performance such as variations in traffic load at specified times
- traffic coming from or going to unexpected locations
- connections made at unusual times
- repeated, failed connection attempts
- unauthorized scans and probes
- non-standard or malformed packets (protocol violations)

If you permit third party (vendor, contractor, supplier, partner, customer, etc.) access to your systems and networks, you must monitor their access to ensure all their actions are authentic and authorized. This includes monitoring and inspecting their network activities.

Why this is important

Monitoring messages as they traverse your network provides the capability to identify intrusive activity at the time it is occurring or soon after. By catching suspicious activity as early as possible, you can immediately begin to investigate the activity and hopefully minimize and contain any damage.

Logs of network traffic may contain evidence of unusual, suspicious, or unexpected activities, indicating that someone has compromised or tried to compromise a system on your network. By inspecting log files on a regular basis, you may be able to identify intruder reconnaissance in advance of an intrusion.

-
1. Internet control message protocol
 2. simple network management protocol

You may also identify attempted or successful intrusions soon after they occur. However, if an intruder has altered log files, this data may no longer be present.

How to do it

- *Notify users³ that network monitoring is being done.*

Inform authorized users of your systems about the scope and kinds of monitoring you will be doing and the consequences of unauthorized behavior.

A common method for accomplishing this is the presentation of a banner message immediately before user login. Refer to the **Other information** section at the end of this practice.

Without the presentation of a banner message or other warning, you can likely not use log files and other collected data in any action you may choose to take against a user.

- *Review and investigate notifications from network-specific alert mechanisms (such as email, voice mail, or pager messages).*

This includes notifications from

- users and other administrators via email or in person
- operating system alert mechanisms
- network and system management software traps such as those that can be set via SNMP (simple network management protocol)
- intrusion detection systems
- custom alert mechanisms from service or application programs (including tools)

- *Review and investigate network error reports.*

These types of notifications typically are produced by

- operating system error reporting mechanisms
- log file filtering tools
- vendor or custom-developed management software
- custom error reporting mechanisms from service or application programs (including tools)

Often an administrator will be able to configure error reporting at a number of criticality, severity, or priority levels when installing the network system, service and application programs, and supporting tools.

- *Review network performance statistics and investigate anything that appears anomalous.*

Statistics are generally produced by vendor or custom performance monitoring tools. Typical statistics include⁴

- total traffic load in and out over time (packet, byte, and connection counts) and by event (such as new product or service release)

-
3. Users are those who access, administer, and manage your systems and have authorized accounts on your systems.
 4. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior.”

- traffic load (percentage of packets, bytes, connections) in and out over time sorted by protocol, source address, destination address, other packet header data
- error counts on all network interfaces
- comparison of previous network performance statistics with current statistics for the same timeframe

Look for

- unexpected changes in performance between current and previously captured statistics; for example network traffic being unusually high or low when compared to expected levels for the time of day and day of the week
- unexpected differences from authoritative network traffic characterization information⁵, such as
 - traffic is coming from unexpected source addresses or using unexpected ports or protocols
 - traffic is going to unexpected destination addresses or using unexpected ports or protocols
 - traffic volume is excessively high or low for time of day and day of week
- unexpected loss of connectivity
- unusual modem activity or availability. Unusual modem activity can indicate intruder access through overlooked entry points (ports) or intruder use of daemon dialers.

➤ *Identify any unexpected, unusual, or suspicious network traffic and the possible implications.*

From network log files and other network traffic collection mechanisms, look for

- reconnaissance (probes, scans, use of mapping tools) in advance of an attack. These can indicate attempts to identify your configuration (hosts, operating systems, network topology, externally accessible paths into your systems, etc.) and identify the Internet service providers (ISP) that you use, along with their configuration.
- connections to or from unusual locations. For example, if a server host is dedicated to a single service (such as serving a public web site), any requests it makes for outbound connections are suspicious. Such requests may indicate that an intruder has compromised the server and that it is being used to launch an attack on another host.
- protocol violations. These include, but are not limited to, invalid option bits in a TCP⁶ packet, invalid sequence numbers in a TCP packet, invalid flags in a TCP packet (ACK before SYN), and invalid fragments. There is no good reason to violate the IP⁷, TCP, ICMP, and UDP⁸ specifications. These types of protocol violations are often a result of an intruder using a network scanner in an attempt to bypass your firewall (that may just check for an established bit set on a packet) and to identify the type of systems on your networks (since different host IP stacks will respond to the error in different ways).

5. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior,” specifically the step “Characterize typical network traffic and performance.”

6. transmission control protocol

7. Internet protocol

8. user datagram protocol

A denial-of-service condition that can occur, for example, when an intruder's host creates TCP half-open connections by sending a flood of SYN packets with no corresponding ACK packets.⁹

- packets with source and destination addresses external to your network. Your firewall should always be configured to prevent this. If it occurs, it may indicate that an intruder has bypassed the firewall, possibly by compromising the firewall host, and is routing their traffic through your network, possibly to take advantage of a network-level trust relationship. It may also indicate the presence of an inside intruder.
- packets with an internal source address that are actually originating from an external source. This can indicate an IP spoofing attack that may have bypassed your firewall.
- unusual port combinations in TCP and UDP packets. This type of traffic could indicate an unexpected service running on the network (such as a backdoor program). It could also indicate that the intruder has bypassed your firewall. Packets with the same source address and a sequence of destination ports often indicate that an intruder is trying to discover both the firewall policy and what services are available on your systems.
- unusual ARP¹⁰ traffic. In a switched network, an intruder can alter the ARP cache on one or more hosts so that any host on the same segment can see traffic on that segment (similar to a network interface card in promiscuous mode on a shared Ethernet segment). The intruder can then gain access to passwords and other unencrypted information sent over the network.
- unusual DHCP/BOOTP¹¹ traffic. An intruder can cause a host to send bogus DHCP replies and convince other hosts that it is their default gateway. The compromised host will then receive all of the traffic for outbound networks and gain access to unencrypted information sent over the network.
- packets with unusual protocol or port numbers sent to broadcast addresses. This type of traffic can indicate a denial-of-service attack.
- an unusually high number of ICMP port unreachable packets from a single host. This indicates that an intruder is scanning the host looking for available services.
- connections made at unusual times
- unusual use of Internet Relay Chat (IRC), a common means of communication used by intruders

- *If you are reviewing network traffic on a system other than the one being monitored, ensure that the connection between them is secure.*

Refer to the module *Securing Network Servers* [Allen 00], specifically the practice "Configure computers for secure remote administration."

Policy considerations

Your organization's networked systems security policy should

- require that users be notified that you will monitor network activities.
- state your objectives for monitoring.
- specify which data streams will be monitored and for what purposes.

9. Refer to CERT advisory CA-96.21, TCP Syn Flooding and IP Spoofing Attacks, available at <http://www.cert.org/advisories>.

10. address resolution protocol

11. dynamic host configuration protocol/boot protocol

- specify the responsibilities and authority of system administrators for handling notifications generated by monitoring and logging software.
 - specify what forms of unexpected network behavior users should watch for. Require users to report any such behavior to their designated security officials and system administrators.
-

Other information

1. For further information on setting up monitoring banners for Windows NT, refer to the implementation “Setting up a logon banner on Windows NT 4.0,” available at <http://www.cert.org/security-improvement/implementations/i034.01.html>. One example of banner language taken from this implementation is

This system is for the use of authorized users only. Individuals using this computer system without authority, or in excess of their authority, are subject to having all of their activities on this system monitored and recorded by system personnel. In the course of monitoring individuals improperly using this system, or in the course of system maintenance, the activities of authorized users may also be monitored. Anyone using this system expressly consents to such monitoring and is advised that if such monitoring reveals possible evidence of criminal activity, system personnel may provide the evidence of such monitoring to law enforcement officials.

2. For further UNIX- and NT-specific network monitoring and network data collection guidance, refer to CERT advisories, incident notes, vulnerability notes, and tech tips, available at <http://www.cert.org>, including
 - Intruder Detection Checklist (http://www.cert.org/tech_tips/intruder_detection_checklist.html)
 - Steps for Recovering from a UNIX or NT System Compromise (http://www.cert.org/tech_tips/win-UNIX-system_compromise.html)
 3. One list of network monitoring tools is contained in the implementation “Identifying tools that aid in detecting signs of intrusion,” available at <http://www.cert.org/security-improvement/implementations/i042.07.html>. Many of these tools can be downloaded from the Center for Education, Research, and Information Assurance Security [CERIAS] (formerly known as Computer Operations, Audit, and Security [COAST]), at <http://www.cerias.purdue.edu/>.
 4. When possible, analyze and correlate data collected from multiple sources (as described in the other practices of this module). Doing some level of correlation analysis (determining when suspicious activity occurring in one part of your infrastructure may be related to suspicious activity in another part) during the intrusion detection process will assist you in determining the full extent of any compromise and its characteristics.¹²
-

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL

<http://www.cert.org/security-improvement/practices/p094.html>

12. Refer to the practice “Analyze all available information to characterize an intrusion” in the module *Responding to Intrusions* [Kossakowski 99].

6

Monitor and inspect system activities for unexpected behavior.

System activities include those associated with system performance, processes, and users¹.

Programs executing on your networked systems typically include a variety of operating system and network services, user-initiated programs, and special-purpose applications such as database services. Every program executing on a system is represented by one or more processes. Each process executes with specific privileges that govern what system resources, programs, and data files it can access, and what it is permitted to do with them.

The execution behavior of a process is represented by the operations it performs while running, the manner in which those operations execute, and the system resources it uses while executing. Operations include computations, transactions with files, devices, and other processes, and communications with processes on other systems via your network.

User activities include login/logout, authentication and other identification transactions, the processes they execute, and the files they access.

If you permit third party (vendor, contractor, supplier, partner, customer, etc.) access to your systems and networks, you must monitor their access to ensure all their actions are authentic and authorized. This includes monitoring and inspecting their system activities.

Why this is important

You need to verify that your systems are behaving as expected and that the processes executing on your systems are attributed only to authorized activities of users, administrators, and system functions.

Unexpected or anomalous system performance may indicate that an intruder is using the system covertly for unauthorized purposes. They may be attempting to attack other systems within (or external to) your network, or they may be running network sniffer programs.

A process that exhibits unexpected behavior may indicate that an intrusion has occurred. Intruders may have disrupted the execution of a program or service, causing it to fail, or to operate in a way other than the user or administrator intended. For example, if intruders successfully disrupt the execution of access-control processes running on a firewall system, they may access your organization's internal network in ways that would normally be blocked by the firewall.

1. Users are those who access, administer, and manage your systems and have authorized accounts on your systems.

How to do it

➤ *Notify users that monitoring of process and user activities is being done.*

Inform authorized users of your systems about the scope and kinds of monitoring you will be doing and the consequences of unauthorized behavior.

A common method for accomplishing this is to present a banner message immediately before user login.²

Without the presentation of a banner message or other warning, you can likely not use log files and other collected data in any action you may choose to take against a user.

➤ *Review and investigate notifications from system-specific alert mechanisms (such as email, voice mail, or pager messages).*

This includes notifications from

- users and other administrators via email or in person
- operating system alert mechanisms
- system management software traps
- intrusion detection systems
- custom alert mechanisms from service or application programs (including tools)

➤ *Review and investigate system error reports.*

These types of notifications typically are produced by

- operating system error reporting mechanisms
- log file filtering tools
- vendor or custom-developed management software
- custom error reporting mechanisms from service or application programs (including tools)

Often an administrator will be able to configure error reporting at a number of criticality, severity, or priority levels when installing the system, service and application programs, and supporting tools.

➤ *Review system performance statistics and investigate anything that appears anomalous.*

Statistics are generally produced by vendor or custom performance monitoring tools.

Typical statistics include³

- total resource use over time (CPU, memory [used, free], disk [used, free])
- status reported by systems and hardware devices such as print queues
- changes in system status, including shutdowns and restarts
- file system status (where mounted, free space by partition, open files, biggest file) over time and at specific times
- file system warnings (low freespace, too many open files, file exceeding allocated size)

2. Refer to the practice “Monitor and inspect network activities for unexpected behavior,” specifically the section titled **Other information**.

3. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior.”

- disk counters (input/output, queue lengths) over time and at specific times
- hardware availability (modems, network interface cards, memory)
- performance statistics meaningful for a specific server or host⁴
- comparison of previous system performance statistics with current statistics

Unexpected shutdowns, reboots, and restarts can indicate the presence of a Trojan horse program that requires a shutdown or restart of a system or service.

► *Continuously monitor process activity (to the extent that you can).*

The examination of processes is complex, time consuming, and resource intensive. The degree to which you are able to identify suspicious processes depends on your knowledge of what processes you normally expect to be executing on a given system and how they should behave.

Due to the large number of processes and their rapidly changing natures, it is impractical for you to monitor them continually yourself. In addition, the amount and value of information that you can gather from a snapshot of currently executing processes may be very limited. This means that you must employ a variety of information-gathering and monitoring mechanisms to help you collect and analyze data associated with processes, and to alert you to suspicious activity.

One common approach with multiuser systems is to set up consoles (or separate terminal windows on workstations) that display the current status of processes and are updated at short intervals. Ideally, these consoles should be hard-wired to the systems for which they are displaying information. With strategic placement of these displays, you can take advantage of the experience of system administrators to notice unexpected activity that may not be picked up by your more immediate alert mechanisms.

► *Identify any unexpected, unusual, or suspicious process behavior and the possible implications.*

As a general guideline, you should look for⁵

- missing processes
- extra processes
- unusual process behavior or resource utilization
- processes that have unusual user identification associated with them

Data from log files and other data collection mechanisms⁶ will help you to analyze the process behavior. These include

- user executing the process
- process start-up time, arguments, file names
- process exit status, time duration, resources consumed

-
4. For example, for a web server, these statistics include pages accessed, connection statistics, user requests over time, which pages are most requested, and who is requesting the pages.
 5. Refer to the implementation “Process analysis checklist” available at [http://www.cert.org/
security-improvement/implementations/i005.02.html](http://www.cert.org/security-improvement/implementations/i005.02.html).
 6. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior.”

- amount of resources used (CPU, memory, disk, time) by specific processes over time; top “x” resource-consuming processes
- system and user processes and services executing at any given time
- the means by which each process is normally initiated (administrator, other users, other programs or processes), with what authorization and privileges
- devices used by specific processes
- files currently open by specific processes

Look for

- processes running at unexpected times
- processes terminating prematurely
- processes consuming excessive resources (wall clock time, CPU time, memory, disk) may warn you of an impending denial-of-service condition or the use of a network sniffer
- unusual processes, such as password cracking, network packet sniffing or any other process not due to normal, authorized activities
- processes with unusually formatted output or arguments (for example, on UNIX systems, a process running as “./telnetd” instead of “/usr/sbin/telnetd”)
- new, unexpected, or previously disabled processes or services. These can indicate that an intruder has installed their own version of a process or service or, for example, are running IRC services, web services, FTP services, and so forth to allow them to distribute tools and files they have stolen (such as password files) to other compromised hosts.
- inactive user accounts that are spawning processes and using CPU resources
- a terminal exhibiting abnormal input/output behavior
- processes without a controlling terminal that are executing unusual programs
- an unusually large number of processes

Pay close attention to the processes associated with intrusion detection and other security tools. Intruders regularly compromise these tools to gain greater leverage and information and to generate decoy alerts to distract and waste the time of system administrators.

➤ *Identify any unexpected, unusual, or suspicious user behavior and the possible implications.*

Data from log files and other data collection mechanisms⁷ will help you to analyze user behavior. These include

- login/logout information (location, time): successful, failed attempts, attempted logins to privileged accounts
- login/logout information on remote access servers that appears in modem logs
- changes in user identify
- changes in authentication status, such as enabling privileges
- failed attempts to access restricted information (such as password files)

7. Refer to the practice “Identify data that characterize systems and aid in detecting signs of suspicious behavior.”

- keystroke monitoring logs
- violations of user quotas

Look for

- repeated failed login attempts including to privileged accounts
- logins from unusual locations or at unusual times including unusual or unauthorized attempts to login via a remote access server
- unusual attempts to change user identify
- unusual processes run by users
- unusual file accesses, including unauthorized attempts to access restricted files
- users logged in for an abnormal length of time (both short and long)
- a user executing an unexpected command
- a user working from an unusual terminal

If you notice unusual activity associated with particular users, initiate supplemental data collection mechanisms to gather detailed information about their activities. Many multiuser systems provide mechanisms to audit all processes associated with a particular user. Since process accounting logs tend to generate a great deal of information rapidly, you will need to allocate sufficient resources to store the data collected. Similarly, detailed network logging of all activity associated with all the systems accessed by a specific user can be voluminous, and you will need to allocate resources accordingly. Review the newly collected data often (at least daily) and rotate files regularly to minimize the amount of information that you have to analyze at any given time.⁸

➤ *Identify other unexpected, unusual, or suspicious behavior and the possible implications.*

If your network interface card is in promiscuous mode, an intruder may be using this mode to run network sniffers for capturing passwords and other sensitive information. Refer to the **Other information** section at the end of this practice. However, keep in mind that legitimate network monitors and protocol analyzers will set a network interface in promiscuous mode as well.

Doing some level of correlation analysis (determining when intrusion activity occurring in one part of your systems may be related to activity in another part) during the intrusion detection process will assist you in determining the full extent of any compromise and its characteristics.⁹

Logging information produced by vulnerability patches (updated software that corrects or closes a vulnerability), if provided by the vendor and if turned on, can aid in identifying a pattern where an intruder exploits more than one vulnerability before gaining access. For example, a failed logged attempt to probe for an old vulnerability (produced by the vulnerability patch) could be followed by a successful probe for a new vulnerability that is not logged. The presence of the vulnerability patch logging information along with other mechanisms such as integrity checking could alert you to this type of intruder action.

8. Refer to the practice “Manage logging and other data collection mechanisms.”

9. Refer to the practice “Analyze all available information to characterize an intrusion” in the module *Responding to Intrusions* [Kossakowski 99].

- *Periodically execute network mapping and scanning tools to understand what intruders who use such tools can learn about your networks and systems.*

We recommend running mapping and scanning tools during non-business hours and when you are physically present because mapping tools can sometimes affect systems in unexpected ways.

Eliminate or make invisible (if possible) any aspect of your network topology and system characteristics that you do not want to be known by intruders who use mapping tools.

- *Periodically execute vulnerability scanning tools on all systems to check for the presence of known vulnerabilities.*

We recommend running such tools during non-business hours and when you are physically present because scanning tools can sometimes affect systems in unexpected ways.

Eliminate all vulnerabilities identified by these tools wherever possible. Many of these can be dealt with by updating configuration file settings and installing vendor-provided patches.¹⁰

Consider using scanning tools that include password analysis as part of the vulnerability assessment. Such analysis may include the identification of weak, non-existent, or otherwise flawed passwords such as those that can be determined using brute force or dictionary-based attacks.

Refer to CERT/CC vulnerability notes¹¹ and “How To Eliminate The Ten Most Critical Internet Security Threats: The Experts’ Consensus, Version 1.25” [SANS 00] for a description of some of the more prevalent vulnerabilities.

- *If you are reviewing system activities on a host other than the one being monitored, ensure that the connection between them is secure.*

Refer to the module *Securing Network Servers* [Allen 00], specifically the practice “Configure computers for secure remote administration.”

Policy considerations

Your organization’s networked systems security policy should

- require that users be notified that monitoring of process and user activities will be done and the objective of such monitoring
- specify the responsibilities and authority of designated systems administrators and security personnel to examine systems, processes, and user activity for unexpected behavior
- specify what forms of unexpected behavior users should watch for. Require users to report any such behavior to their designated security officials and system administrators.
- specify what software and data users and administrators are permitted to install, collect, and use, with explicit procedures and conditions for doing so
- specify what programs users and administrators are permitted to execute and under which conditions

10. Refer to the module *Securing Network Servers* [Allen 00], specifically the practice “Keep operating systems and applications software up to date.”

11. Available at http://www.cert.org/vul_notes

Other information	<ol style="list-style-type: none"> 1. One common activity of intruders is to gather information from the traffic on your networks to find user account names, passwords, and other information that may facilitate their ability to gain access to your systems. They do this by breaking into one system on your network and installing and executing a sniffer program. This program collects information about connections established between systems from network data packets as they arrive at or pass by the compromised system. To hide this illicit activity on compromised systems, intruders typically modify log files and replace programs that would reveal the presence of the sniffer program with Trojan horse versions. The substitute programs appear to perform the same functions but exclude information associated with the intruders and their activities. In many documented cases of this type of intrusion, the intruders' activities went unnoticed for a considerable amount of time, during which they collected enough information to gain privileged access to several other systems. <p>This underscores the importance of using verified software to examine your systems¹² and the need to verify the integrity of your files.¹³ Unfortunately, there are several sophisticated collections of programs that intruders can use to rapidly gain access to systems and "set up shop" to install and execute a sniffer. This means that the only method you may have to catch such activity is to use verified software to examine processes on your systems for unexpected behavior.¹⁴</p> <p>Processes associated with a sniffer will typically have transactions with a network interface that has been placed in <i>promiscuous mode</i>¹⁵, as well as a file or network connection to which the information gathered from network packets is being sent.</p> <p>Refer to CERT Advisory CA-94.01, Ongoing Network Monitoring Attacks, available at http://www.cert.org/advisories.</p> <ol style="list-style-type: none"> 2. One list of system monitoring tools is contained in the implementation "Identifying tools that aid in detecting signs of intrusion," available at http://www.cert.org/security-improvement/implementations/i042.07.html. Many of these tools can be downloaded from the Center for Education, Research, and Information Assurance Security [CERIAS] (formerly known as Computer Operations, Audit, and Security [COAST]), at http://www.cerias.purdue.edu/. 3. When possible, analyze and correlate data collected from multiple sources (as described in the other practices of this module). Doing some level of correlation analysis (determining when suspicious activity occurring in one part of your infrastructure may be related to suspicious activity in another part) during the intrusion detection process will assist you in determining the full extent of any compromise and its characteristics.¹⁶ <hr/> <ol style="list-style-type: none"> 12. Refer to the practice "Ensure that the software used to examine systems has not been compromised." 13. Refer to the practice "Inspect files and directories for unexpected changes." 14. Refer to the practice "Monitor and inspect system activities for unexpected behavior." 15. Network interfaces on most systems normally operate in <i>non-promiscuous mode</i>, which means that they ignore network packets not explicitly addressed to them. In <i>promiscuous mode</i>, no packets are ignored, that is, all packets that traverse the network segment to which the system is attached are read by its network interface and are accessible to processes executing on that system. 16. Refer to the practice "Analyze all available information to characterize an intrusion" in the module <i>Responding to Intrusions</i> [Kossakowski 99].
--------------------------	--

Where to find updates	The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL http://www.cert.org/security-improvement/practices/p095.html
------------------------------	---

7

Inspect files and directories for unexpected changes.

The file systems in your network environment contain a variety of software and data files. Unexpected changes in directories and files, especially those to which access is normally restricted, may indicate an intrusion. Changes could include modifying, creating, or deleting directories and files. What makes such changes *unexpected* may depend on who changed them, and where, when, and how the changes were made.

If you permit third party (vendor, contractor, supplier, partner, customer, etc.) access to your systems and networks, it is critical that you actively monitor their access to your systems and networks as well as any processing they do. This helps ensure that all actions are authentic and authorized. Monitoring access includes examining all relevant directories and files.

Why this is important

Intruders often create, substitute, modify, and damage files on systems to which they have gained access. To hide their presence on your systems, it is common for intruders to replace system programs with substitutes that perform the same functions but exclude information that could reveal their illicit activities. They also often modify system log files to remove traces of their activities¹. By masking their presence on a compromised system, intruders prolong the time they have to use that system for their purposes. In several notable cases, the presence of intruders on compromised systems was not discovered until many months after the initial intrusion occurred.

Intruders may also create new files on your systems. For example, they may install backdoor programs or tools used to gain privileged access on the system. Intruders also make use of the disk space on compromised systems to store their tools and other artifacts.

Private data files and files containing mission-critical information are common targets of modification or corruption by intruders. Information about your organization that is accessible to the public or to subscribers via public networks and the Internet is also a common target. Several documented cases exist of prominent organizations that have had their web sites modified to include offensive content and other erroneous information.

How to do it

► *Establish priorities and schedules.*

Examine the directories and files on your system and prioritize how frequently you should check them. The more mission- or security-critical the file, the more frequently you should check it.

1. Intruder actions of this type can be accomplished using tools that are part of rootkit. For Windows NT, refer to <http://www.rootkit.com>. For Linux, refer to <http://www.securityfocus.com/tools/1489>. To check for signs of the presence of rootkit, refer to <http://www.securityfocus.com/tools/1646>.

We recommend that checking be done at least daily, preferably at the start of the business day, to cover all processing done during the immediately preceding 24 hours.

► *Verify the integrity of directories and files according to your established schedule.*

Compare the attributes and contents of files and directories to the authoritative reference (either complete copies or cryptographic checksums).² Identify any files and directories whose contents or other attributes have changed.

Always access authoritative reference data directly from its secured, read-only media. Never transmit authoritative reference data over unsecured network connections unless you use mechanisms such as digital signatures and cryptographic checksums to verify data integrity.

► *Identify any unexpected, unusual, or suspicious changes to files and directories and their possible implications.*

Data from log files and other data collection mechanisms will help you to analyze changes to files and directories.³ These include

- cryptographic checksums for all files and directories
- lists of files, directories, attributes
- accesses (open, create, modify, execute, delete), time, date
- changes to sizes, contents, protections, types, locations
- additions and deletions of files and directories
- results of virus scanners

Look for

- unexpected file or directory access, creation, or deletion
- unexpected changes to file or directory protections or access control lists. Identifying these can aid, for example, in detecting the creation of files in user home directories that can be later used for backdoor access. Improperly set access control lists on system tools may indicate that an intruder has located and executed security tools that were installed by the authorized system administrator.
- unexpected changes to file or directory sizes, contents, and other attributes. These may signify that a file or service has been replaced with the intruder's version including the installation of a Trojan horse or backdoor. An intruder inadvertently enabling debugging can easily quadruple the size of a file.
- unexpected changes to password files such as unauthorized creation of new accounts and accounts with no passwords
- unexpected changes to system configuration files and other restricted and sensitive information including firewall filtering rules
- unusual or unexpected open files. These can reveal the presence of sniffer logs or programs.

2. Refer to the practice "Identify data that characterize systems and aid in detecting signs of suspicious behavior."

3. Ibid.

- violations of log file consistency (unexpected changes in file size, gaps in time between log records)
- presence of viruses, backdoors, and Trojan horses detected by scanning tools⁴

Compromised systems that support a promiscuous network interface can be used by intruders to collect host and user authentication information that is visible on the network. Sniffers are able to capture user keystrokes containing host, account, and password information. The presence of sniffers can be detected by looking for Trojan horse programs, suspect processes, unexpected ASCII files, and modifications to files.⁵

Policy considerations	Your organization's networked systems security policy should <ul style="list-style-type: none">• require that users be notified that file and directory examination will be done and the objective of such examination• specify the responsibilities and authority of designated systems administrators and security personnel to examine files and directories on a regular basis for unexpected changes• require users to report any unexpected changes to their software and data files to system administrators or your organization's designated security point of contact
Other information	<ol style="list-style-type: none">1. Some kinds of important files, such as log files and database tables, are expected to change frequently (perhaps several times per second). In general, the techniques described above will not be useful in distinguishing normal changes to these file types from those that might have been caused by intruders. Techniques based on transaction auditing are more useful in these cases.2. When possible, analyze and correlate data collected from multiple sources (as described in the other practices of this module). Doing some level of correlation analysis (determining when suspicious activity occurring in one part of your infrastructure may be related to suspicious activity in another part) during the intrusion detection process will assist you in determining the full extent of any compromise and its characteristics.⁶

Where to find updates	The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL http://www.cert.org/security-improvement/practices/p096.html
------------------------------	---

-
4. Refer to the module *Securing Network Servers* [Allen 00], specifically the practice "Protect computers from viruses and similar programmed threats."
 5. Refer to the practice "Monitor and inspect system activities for unexpected behavior," specifically **Other information**.
 6. Refer to the practice "Analyze all available information to characterize an intrusion" in the module *Responding to Intrusions* [Kossakowski 99].

8

Investigate unauthorized hardware attached to your organization's network.

Unauthorized hardware may include computers connected to network segments or hubs, and peripheral communication or input/output equipment such as modems, terminals, printers, and disk or tape drives.

Why this is important

Intruders actively attempt to circumvent network perimeter defenses. If they can gain physical access to your organization's internal network, then they can install their own equipment and software. Alternatively, intruders may learn of insecure (unauthorized) equipment added by users that they can use to gain access to your organization's network. For example, users might install modems for the purpose of remote access to their office computers from home. Intruders often use automated tools to identify modems attached to public telephone lines. If the configuration of the dial-up access, and the traffic through it, is not secured, intruders may use such backdoors to gain access to the internal network, bypassing preventive measures that may have been put in place to restrict external connections to your organization's network. They may then capture network traffic, infiltrate other systems, disrupt operations, and steal sensitive, private information.

Access to other peripheral equipment may also facilitate intrusions. Unsecured output and removable media devices, such as printers and disk drives, may give intruders the opportunity to generate copies of sensitive information that can be physically removed from your organization's premises.

How to do it

- *Perform a monthly "walkabout" audit of all systems and peripherals attached to the network infrastructure.*

Visits to physically examine equipment attached to the network should not be announced, so that unauthorized equipment cannot be hidden before the auditors arrive.

Using your documented hardware inventory¹, identify any missing hardware, hardware that is not in its expected location, and any unexpected, extra hardware.

- *On a daily basis, probe for unauthorized modems attached to your organization's telephone lines.*

You can do this using *daemon dialer* tools. Because this process causes all dialed telephones to ring, we recommend that it be done outside of normal working hours.

1. Refer to the practice "Identify data that characterize systems and aid in detecting signs of suspicious behavior," specifically the step "Generate an inventory of your system hardware."

- *On a daily basis, probe all internal network segments to identify unauthorized hardware attached to your network.*

Examine

- unauthorized devices attached to your network
- any new or unexpected IP or MAC addresses
- any new or unexpected network ports on switches

You can do this using public domain tools such as ARPWATCH² and a variety of commercial network management software packages.

Identify any missing hardware, hardware that is not in its expected location, and any unexpected, extra hardware.

- *On a daily basis, look for unexpected routes between the organization's network and external networks.*

Examine the network traffic logs for connections that originate outside your network and are destined for addresses outside your network. Traffic that transits in this way could indicate that an unauthorized computer is connecting to one of your hosts.

If possible, compare the network traffic logs from individual hosts/workstations with network traffic logs from the firewall host(s). Discrepancies or mismatches could indicate that traffic is being routed through unsecured connections or gateways directly to the individual host, bypassing your organization's firewalled Internet connection.

Policy considerations	Your organization's networked systems security policy should <ul style="list-style-type: none">• require the maintenance of documented hardware inventories• require the maintenance of a documented network topology• specify the authority and responsibility of designated security personnel to<ul style="list-style-type: none">– perform physical audits of installed hardware and software– establish network connections and routes• specify what kinds of hardware and software users are permitted to install on their desktop workstations
Other information	In addition to the periodic inspections of hardware recommended above, you may need to conduct inspections in response to suspected intrusions. Watch for evidence of activities that indicate unusual access to your network. ³
Where to find updates	The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL http://www.cert.org/security-improvement/practices/p097.html

2. Available at <ftp://ftp.ee.lbl.gov/>. ARPWATCH is only effective for hosts attached to your local area network as external hosts are represented by your router/firewall.

3. Refer to the practice "Monitor and inspect network activities for unexpected behavior."

9

Inspect physical resources for signs of unauthorized access.

Although we tend to think of the information in networked computer systems as being in electronic form, we should remember that this information is held on physical media—CD-ROMs, tapes, disks, paper—and as objects are subject to physical compromise by theft, destruction, corruption, or unauthorized duplication. To ensure the security of your network, you should also ensure the physical security of its components by periodically inspecting them for possible compromise.

In many organizations, designated personnel are responsible for the physical security of the premises. However, as a system or network administrator, you are often in a unique position to notice signs of physical access to system resources.

Why this is important

If a document or electronic storage medium is stolen, the confidentiality and availability of the information it contains is lost. Even if the item is recovered, you won't know the extent to which its contents have been copied and disseminated. Also, you won't know whether the information it contains has been corrupted or altered. Furthermore, if the compromised information is critical to security, (e.g., user passwords, internal network addresses, or system configuration data) your entire network is potentially threatened by more damaging intrusions.

Therefore, it is just as important for you to keep track of physical resources and to promptly detect attempts at physical intrusion and access as it is for you to track and protect your electronic resources.

How to do it

- *Daily, check all physical means of entrance or exit for signs of tampering, trespassing, or attempted trespassing.*

Keep in mind that intruders have many strategies for obtaining confidential or security-critical documents. For example, they may steal discarded copies of reports, console logs, system printouts, or other sensitive data. They search through trash containers or dumpsters to find carelessly discarded physical copies. They may also attempt to steal backup or archive tapes, whose disappearance may not be noticed for some time.

- *Daily, check physical resources for signs of tampering.*

For example, inspect locks or seals on hardware cabinets, review console logs, and monitor paper usage.

- *Regularly perform a physical audit of all movable media.*

We recommend that this be done weekly, if possible.

Ensure that write-disabled media continue to be so.

Note that, as a complementary practice, you should also audit the contents of the media for electronic integrity.

- *Report all signs of unauthorized physical access to your organization's internal security point of contact.*

This includes access to offsite data storage and disaster recovery sites.

Policy considerations	Your organization's networked systems security policy should <ul style="list-style-type: none">• require the tagging and inventory of all physical computing resources¹• specify how to respond when a physical intrusion has been detected
Other information	You may want to consider encrypting all backup and other selected electronic media in the event that your site, an offsite data storage site, or a disaster recovery site is physically compromised.
Where to find updates	The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL http://www.cert.org/security-improvement/practices/p098.html

1. Refer to the practice "Identify data that characterize your systems and aid in detecting signs of suspicious behavior," specifically the step "Generate an inventory of your system hardware."

10

Review reports by users and external contacts about suspicious and unexpected behavior.

In security-conscious organizations, users¹ will report suspicious events and behaviors. You, as a system or network administrator, should use those reports along with information you gather, to help identify possible intrusions. When appropriate, you should also use external sources of information, such as reports from incident response teams, to help you in deciding whether or not you need to augment your monitoring and incident analysis efforts.²

Why this is important

Recruiting users and external contacts to assist you in security monitoring greatly extends your ability to detect intrusions, potentially allowing you to detect intrusions of which you were previously unaware. Not only do they add to the number of people alert to possible intrusions, they can often be more aware of the “normal” behavior of their personal computing environments than you are. Many intrusions are not discovered until someone with day-to-day experience using a particular system notices something unusual. Users are susceptible to intruder-initiated social engineering attempts (for example to obtain passwords or to gain physical access) and need to understand how to identify and report these.

Intruders often compromise multiple systems when they attack a target site. At each compromised system, there may be telltale signs of intrusive activities that users of the system discover. Although a single user report may not be sufficient evidence of an intrusion, analysis of several reports may reveal a pattern of attack underway. By consolidating users’ reports of suspicious system behaviors, you may also be able to determine how widespread attacks against your networked systems are.

Administrators from other organizations may contact you if they have a reason to believe that an intrusion into their systems may involve or affect your organization. Always thoroughly investigate any reports you receive from incident response teams, such as the CERT/CC, to determine if, in fact, an intrusion has occurred at your site. If your network environment supports connections to external networks, it is possible that your systems may have been compromised and are serving as unwitting participants in a large-scale attack (such as a distributed denial-of-service attack) against several sites.

-
1. Users are those who access, administer, and manage your systems and have authorized accounts on your systems.
 2. Refer to **Other information** in the practice “Establish a policy and procedures that prepare your organization to detect signs of intrusion.”

How to do it	<ul style="list-style-type: none"> ➤ <i>Perform “triage” upon receipt of a report.</i> <p>Immediately gather as much information as necessary to make an initial assessment of whether there is a probable intrusion and, if so, how severe it seems to be. You may need to make direct contact with the user to get a description of what was observed. Also acquire any records or data from logging, monitoring, or other data collection mechanisms that illustrate the problem. If the information clearly indicates an intrusion attempt, investigate it immediately.</p> <ul style="list-style-type: none"> ➤ <i>Evaluate, correlate, and prioritize reports.</i> <p>On a regular basis (daily, if possible), review all user and external reports. These include new reports, reports currently under investigation, and any reports that remain unresolved after investigation. Look for correlations or patterns among the reports. Prioritize and schedule investigations of all reports based on your assessment of their severity. If it is unfounded, close the report and provide feedback to the user who reported the problem.</p> <ul style="list-style-type: none"> ➤ <i>Investigate each report or set of related reports.</i> <p>Based on the nature of the report, you may need to contact other users to document their observations. You may also need to verify the integrity of directories and files³, examine your system and network logs⁴, examine processes on affected systems⁵, and install additional monitoring mechanisms to identify the cause of the anomalous behavior.</p> <ul style="list-style-type: none"> ➤ <i>If an intrusion is detected, initiate your intrusion response procedures immediately.⁶</i> <p>Report suspicious behavior or signs of intrusion to your organization’s designated security point of contact.</p> <ul style="list-style-type: none"> ➤ <i>Document and report your findings.</i> <p>Regardless of your investigation’s outcome, record your findings and report them to the users who submitted the reports, the system and network administrators, the security personnel in your organization, and other appropriate individuals, specified in your organization’s policies.</p>
---------------------	--

Policy considerations	Your organization’s networked systems security policy should <ul style="list-style-type: none"> • require users to report any unexpected or suspicious system behavior immediately to their designated security official and system administrator • require users to report any physical intrusions to networked systems or offline data storage facilities immediately to their designated security official and system administrator
------------------------------	--

-
3. Refer to the practice “Inspect files and directories for unexpected changes.”
 4. Refer to the practices “Monitor and inspect network activities for unexpected behavior” and “Monitor and inspect system activities for unexpected behavior.”
 5. Refer to the practice “Monitor and inspect system activities for unexpected behavior.”
 6. Refer to the module *Responding to Intrusions* [Kossakowski 99].

- require system administrators to investigate each reported suspicious activity to determine whether it represents an intrusion
 - require system administrators to notify users in advance of any changes that will be made to the systems they use, including software configurations, data storage and access, and revised procedures for using systems as a result of the changes
-

Other information

A report should include the following information:⁷

- contact information for the individuals discovering the problem and any responsible parties that are involved (such as the system administrator)
- target systems and networks and all of their characteristics such as operating system versions and IP addresses
- purpose of the systems under attack, including the types of services and applications they provide, as well as an indication of the importance or criticality of the system
- any evidence of intrusion, including method of attacks used, vulnerability exploited, source IP address of attacker, and network contact information for this address
- list of parties to notify, such as legal, other technical, management, and public relations⁸

Refer also to the CERT/CC incident reporting guidelines available at http://www.cert.org/tech_tips/incident_reporting.html.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL

<http://www.cert.org/security-improvement/practices/p99.html>

7. Excerpted from “Do you have an intrusion detection response plan?” [Reavis 99]

8. Refer to the module *Responding to Intrusions* [Kossakowski 99], specifically the practice “Communicate with all parties that need to be made aware of an intrusion and its progress.”

11

Take appropriate actions upon discovering unauthorized, unexpected, or suspicious activity.¹

You need to record and investigate unexpected activity and determine if it should be reflected in your characterization baseline, alerting, or other data collection mechanisms.

If the activity warrants (for example, unauthorized access to sensitive data), you need to initiate your intrusion response procedures.

Why this is important	Identifying unauthorized or suspicious activities and then <i>not</i> taking appropriate follow-up actions will perpetuate any damage or other negative consequences caused. These consequences include possible loss of integrity, availability, or data confidentiality, and legal liability. In addition, these activities are likely to recur, placing your systems at considerable risk in the future.
------------------------------	---

How to do it	➤ <i>Document any unusual behavior or activity that you discover.</i>
---------------------	---

Over time, you may see recurring kinds of unusual or suspicious activity. Maintaining

records of these activities and noting your conclusion of their causes will help you and others to understand new occurrences more quickly and accurately.

For example, in *Network Intrusion Detection* [Northcutt 99], Northcutt writes

To detect and classify a coordinated attack [one coming from or going to multiple locations], it helps to have a database of all traffic and techniques to complement your signatures [of known attacks]. Without a database of traffic that covers a time window of at least a couple months, there is no way to determine whether this activity [that you are now investigating] has been going on and simply hasn't been detected, or whether it is a new pattern. [p. 171]

Northcutt also recommends creating a directory to store data traces. The data traces can be examined when investigating an unknown attack pattern.

➤ <i>Investigate each documented anomaly.</i>

Ask yourself questions such as

- Is the apparent anomaly the result of a legitimate new or updated characteristic of your system? (e.g., the unexpected process is executing a recently added administrative tool)

1. Refer to the other practices in this module for examples of such activity.

- Can the anomaly be explained by the activities of an authorized user? (e.g., the user really was in Cairo last week and connected to the network, a legitimate user made a mistake)
- Can the anomaly be explained by known system activity? (e.g., there was a power outage that caused the system to reboot)
- Can the anomaly be explained by authorized changes to programs? (e.g., the mail log showed abnormal behavior because the system programmer made a mistake when the software was modified)
- Did someone attempt to break into your system and fail?
- Did someone break in successfully? Do you have the data that will tell you what they did?

► *Recognize the iterative nature of analysis and investigation.*

Often you will observe an initial indication of suspicious behavior but will not have sufficient information to determine what occurred. You may need to

- look for past occurrences of similar behavior and study the results of that investigation
- formulate and ask different questions to better identify what data will best reveal what happened
- modify the configuration of selected data collection mechanisms² to collect additional data or better filter and select from existing data
- add new data collection mechanisms

► *If any activity or event cannot be attributed to authorized or explainable activity, initiate your intrusion response procedures immediately.³*

Report such occurrences to your organization's designated security point of contact.

► *Update the configuration of alert mechanisms if required.*

This is warranted if a previous event notification that occurred via logs, error reports, statistics reports, or another data collection mechanism is now of a priority that an alert is called for.

The reverse is also true. An event that is reported as an alert all of the time may become less important and need to be changed to be captured as an error report.

► *Update all characterization information⁴ to reflect what you learn from reviewing any unusual activity or event.*

This is important

- if an unusual activity occurs frequently enough for you to consider it normal and expected such that you should add it to an asset's characterization baseline

2. Refer to the practice "Identify data that characterize systems and aid in detecting signs of suspicious behavior."

3. Refer to the module *Responding to Intrusions* [Kossakowski 99].

4. Refer to the practice "Identify data that characterize systems and aid in detecting signs of suspicious behavior."

- if a new activity has occurred and needs to be added to an asset's characterization baseline
 - if a previously normal or expected activity should now be considered suspicious or unexpected
 - if a previously normal or expected activity should be dropped from consideration for analysis altogether
- *Update all logging and data collection mechanism configurations to reflect information on new attack methods.⁵*
- *Dispose of every reported event.*

This can include

- resolution and closure
- deciding not to pursue the reported event further unless it becomes more critical
- taking no action but preserving the reported event to see if it recurs or contributes to a pattern of events

Policy considerations

Your organization's networked systems security policy should

- specify the actions to be taken following the discovery of unexpected, unusual, or suspicious activity
- require that the actions prescribed are actually performed
- specify the responsibilities and authority of designated systems administrators and security personnel to take the prescribed actions

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the web at URL

<http://www.cert.org/security-improvement/practices/p100.html>

5. Refer to **Other information** in the practice "Establish a policy and procedures that prepare your organization to detect signs of intrusion."

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.			
1. AGENCY USE ONLY (leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	October 2000	Final	
4. TITLE AND SUBTITLE	5. FUNDING NUMBERS		
Detecting Signs of Intrusion	C — F19628-95-C-0003		
6. AUTHOR(S)			
Julia Allen, Ed Stoner			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER		
Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213	CMU/SEI-SIM-009		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER		
HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			
11. SUPPLEMENTARY NOTES			
12.a DISTRIBUTION/AVAILABILITY STATEMENT		12.b DISTRIBUTION CODE	
Unclassified/Unlimited, DTIC, NTIS			
13. ABSTRACT (maximum 200 words)			
This security improvement module, <i>Detecting Signs of Intrusion</i> , describes practices involved in preparing to detect and detecting intrusions into networked computer systems. The practices are designed to help network and system administrators prepare for and detect intrusions by looking for unexpected or suspicious behavior and then recognizing "fingerprints" of known intrusion methods.			
14. SUBJECT TERMS		15. NUMBER OF PAGES	
computer attackers, computer break-ins, computer intrusions, computer security, detecting intrusions, Internet security, intruders, intrusion detecting, network monitoring, network security, security, web servers, system monitoring, World Wide Web		64	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102